EDN®

AUG 2012
www.edn.com
Issue 14

THE FUTURE OF
CONNECTED-DEVICE SECURITY

Page 38

1011

- 9 OCT 2012

BXNDNSK **************AUTO**SCH**5-DIGIT 77079
0002976515/4#
U4 EDN_2 8/23/12 0022   #122 #30356
ACHLAN MAHER - TECHNICAL DIRECTOR
ETECTION SYSTEMS
4781 MEMORIAL DR
OUSTON TX 77079-5210

EDN
contents
August 2012

The future of connected-device security

38 Designers building scalable data protection into next-generation electronics confront a plethora of design constraints and technology choices, building up from the hardware and software roots of trust.
by David Kleidermacher, Green Hills Software

Mixed-signal SOC verification using analog behavioral models

33 Bringing analog and mixed-signal blocks to a higher abstraction level yields more effective mixed-signal simulation and more complete verification environments.
by Qi Wang, Cadence Design Systems

IMAGE: ISTOCK

Managing the 8- to 32-bit processor migration

49 To prevent Lilliputian processor disparities from hatching a swarm of bugs, application developers working in C must keep native processor size and endianness in mind and write their code with an eye toward portability.
by Kevin King, Renesas Electronics America

DESIGNIDEAS

contents August 2012

pulse
Dilbert 14

DEPARTMENTS & COLUMNS

66

70

[www.edn.com]

Easy POL Converter Design

Online Tool Simplifies Sup*IR*buck® POL Design http://mypower.irf.com/Sup*IR*buck

Integrated Voltage Regulator Product Offering

	5V bias required			Integrated LDO	
	1- 16Vin	1-21Vin	3-27Vin	1- 16Vin	1-21Vin
1-3A	IR3843W				
4-6A	IR3842	IR3853		IR3839*	IR3897*
	IR3842W	IR3863			IR3898*
	IR3832W*	IR3856	IR3473		
7-10A	IR3831W*	IR3859		IR3838*	IR3899*
	IR3841	IR3865	IR3475	IR3865	
	IR3841W				
11-16A	IR3840		IR3476	IR3837*	IR3894*
	IR3840W		IR3477		IR3895*
16-35A	iP1827				
	iP1837				

*DDR Tracking

Size(mm)

7.7 x 7.7		12-35A
5 x 6	2-16A	
4 x 5	1-10A	
	Sup*IR*Buck®	*iPOWIR*®

Current (A)   15   35

The Sup*IR*buck® Advantage:

- Best in class efficiency
- Common footprint over wide load range for design flexibility
- Full featured and protected for easy on-board implementation
- Integrated controller, bootstrap diode,control FET and Sync FETKY

The Tool Provides:

- Schematics
- Bill of Materials
- Simulation and Thermal Analysis

for more information call 1.800.981.8699 or visit us at www.irf.com

International IOR Rectifier
THE POWER MANAGEMENT LEADER

EDN online

JOIN THE CONVERSATION
Comments, thoughts, and opinions shared by *EDN*'s community

In response to "Rounding errors, verification, and laziness," a blog post at **www.edn.com/4391628**, UD-00001 commented:

"Sorry to say, but it's not laziness; it tends to be a lack of education.

"During my university days, integer math was taught. Nowadays, it seems that algorithmic precision is no longer an issue: 'Hey, don't think; use float!'
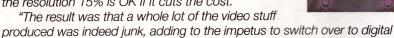
"They are not even aware of the degree of incertitude. And in my current job in automotive electronics, there often is quite a difference between the float results and precisely designed integer math—making the difference between superior and sluggish performance.

"But I understand; integer is not sexy."

In response to "Revisiting the analog video decoder: brushing up on your comb filters," at **www.edn.com/4390178**, WKetel commented:

"The potential for problems with NTSC video always comes back to cost. Saving half a cent but reducing the resolution 15% is OK if it cuts the cost.

"The result was that a whole lot of the video stuff produced was indeed junk, adding to the impetus to switch over to digital video. Nobody really complained about the poor picture quality, because that was what they had been taught to expect. It used to be better, though.

"The big benefit of digital, in obsoleting all of the analog sets, [was in] opening a whole new market. The next step will be some digital encoding, resulting in pay per view for everything.

"Not just yet, but give it another year or two."

EDN invites all of its readers to constructively and creatively comment on our content. You'll find the opportunity to do so at the bottom of each article and blog post.

CONTENT
Can't-miss content on EDN.com

SLIDESHOW: AFTER THE MARS ROVER—A MINING COLONY IN 2039

NASA's most advanced Mars rover, Curiosity, has landed on the Red Planet. The one-ton rover, hanging by ropes from a rocket backpack, touched down on Mars to end a 36-week flight and begin a two-year investigation. What's next? President Obama has laid out a bold vision for sending humans to Mars in the mid-2030s.

www.edn.com/4391812

WHITE LEDS PRINTED ON PAPER: A DOCTORAL THESIS

Here is a doctoral thesis, provided as a multipart article, demonstrating that white LEDs made from zinc oxide and a conducting polymer can be manufactured directly on paper.

www.edn.com/4391796

ENGINEERING COMMUNITY
Opportunities to get involved and show your smarts

designeast
center of the engineering universe

DESIGN EAST

Will embedded programming be dead by 2030? Which disruptive technologies and consumer innovations are on the verge of hitting critical mass? What's the latest in Android, LEDs, and sensors? Explore the hottest developments in engineering, build knowledge and skills, network with peers and technical experts, discover cutting-edge products and technologies, and find solutions to your biggest design challenges at Design East, September 17 to 20 at the Hynes Convention Center in Boston: **http://east.ubmdesign.com**.

BY PATRICK MANNION, BRAND DIRECTOR

Next-gen "Edisons" doing just fine

The notion that engineering's best and brightest are falling back on iPhone and Android apps instead of doing "real" engineering was put to rest in Austin this month at NIWeek's Graphical Systems Design Achievement Awards. I saw this fallacy for what it was as I prepared to announce that Peritec Inc's Naoki Hoshino had won the Editor's Choice honor for his work on a desktop EMI evaluation system.

Naoki Hoshino of Peritec accepts Patrick Mannion's congratulations on his Editor's Choice honor at NIWeek's Graphical Systems Design Achievement Awards.

We've seen the footage from revered colleges, such as Stanford, showing dorm rooms packed with the next generation of engineers seemingly focused more on smartphone apps than on the next tech breakthrough. We've heard OEMs complain that engineers are coming out of college ill prepared for the engineering world. It's easy to despair over what those signs might mean for the future of the tech industry.

But when I attend an event like the NIWeek GSDA awards (which *EDN* parent UBM Electronics co-sponsored), I emerge reinvigorated and brimming with optimism. I meet young innovators who are wrestling with real-world problems, from clinical assessment of conditions such as osteoporosis using ultrasound and lasers to off-grid energy solutions and earthquake monitoring.

The event's engaging master of ceremonies, Dave Wilson, director of academic and training and certification marketing for National Instruments, kept the proceedings moving as he deftly described every paper. What resonated with me most, however, were his closing remarks on technical audaciousness and sheer ingenuity.

For some in the *EDN* community, a design based on LabView combined with a modular PXI or CompactDAQ box may not epitomize "true" design. But I really liked what James Truchard, NI's cofounder, president, and CEO, had to say during his keynote address the day
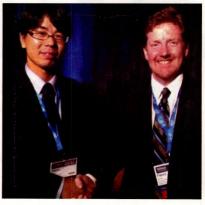
Hoshino shows off his project on the show floor. Turns out he's a regular guy—a really smart one.

before the GSDA event: "Our goal is to make all these Edisons make progress."

I wonder what Thomas Edison himself would have done with LabView, NI hardware, and some Xilinx FPGAs. The possibilities boggle the mind.

For his part, Editor's Choice honoree Hoshino's paper was on "Developing an EMI desktop evaluation system using a new measurement technique" (http://bit.ly/P4xOoK). I chose it as the winner because the problem was technically challenging; because the work has long-term implications; and, in the spirit of the awards, because the achievement itself democratizes innovation by greatly reducing the cost of evaluating a design for EMI by eliminating the need to buy or rent space at an anechoic chamber.

Of course, the work of stellar performers such as Hoshino and the GSDA winners won't quell the calls for "better" education or "more qualified" engineers, but what do such terms even mean anymore? Engineering has changed on many levels, and the emerging generation of engineers reflects those changes. Sometimes, as in Hoshino's case, the innovators themselves change the way engineering is done. The one constant: The work is about solving problems.

I'm curious to know how newer engineers and their more experienced counterparts truly view each other across the generational divide. I'd define a "new engineer" as one who has been out of college for 10 years or less. What do you think? **EDN**

Contact me at patrick.mannion@ubm.com.

Here's the full EMI measurement system in operation. On the left is the PXI box with LabView running on top. To the right, the EMI probe measures the EMI directly over an IC.

pulse

INNOVATIONS & INNOVATORS

LED increases efficacy by 20%

Cree Inc's new XLamp XP-G2 high-brightness LED produces up to 20% more lumens per watt and 2.5 times the lumens per dollar compared with the company's original XP-G LED. The brighter, more efficient XP-G2 LED provides customers with an immediate boost in performance and lets luminaire manufacturers either use fewer LEDs to get the same brightness at lower cost or increase brightness levels using the same LED count and power.

Characterized and binned at 85°C, the XP-G2 LED leverages the same 3.45×3.45-mm footprint as the XP-G LED and is compatible optically with the original part, providing drop-in-ready performance enhancements to shorten the LED-fixture design cycle and improve customer time to market. The LED can enable a broad range of high-lumen applications, from indoor and outdoor installations to portable and lamp retrofits.

Built on Cree's new SC³ technology platform, XP-G2 LEDs combine high light output, reliability, and efficacy to deliver up to 151 lumens/W at 350 mA, 85°C or 165 lumens/W at 350 mA, 25°C in cool white (6000K). In warm white (3000K), the XP-G2 LED delivers up to 133 lumens/W at 350 mA, 85°C or 145 lumens/W at 350 mA, 25°C. The silicon-carbide-based SC³ technology platform features advancements in LED chip architecture and phosphor formulation, as well as a new package design.

Luminaire makers seeking Energy Star qualification will have access to specification and performance data, including LM-80 reports, which can speed time to market. XP-G2 LEDs are a "successor" product to the original XP-G LED for LM-80 data; that status accelerates the qualification of luminaires using just 3000 hours of LM-80 data instead of the normal 6000 hours.

The ROHS- and REACH-compliant XP-G2 LED features a maximum drive current of 1.5A and a level 4 rating and is UL-recognized. Distributor Mouser lists the XP-G2 in cool white at $5.18 in sample quantities.—**by Margery Conner** ▷**Cree**, www.cree.com.

The brighter, more efficient XP-G2 LED provides an immediate boost in performance and enables manufacturers to use fewer LEDs to get the same brightness at lower cost or increase brightness levels using the same LED count and power.

Name
Dr. Dave Barrett

Job Title
*Professor,
Mechanical Engineering*

Area of Expertise
Robotics

LabVIEW Helped Me
*Bridge the gap between
teaching theory and real-
world design experience*

Latest Project
*Building a robotic tuna
to swim across
the Atlantic Ocean*

NI LabVIEW

LabVIEW makes me better by making complex

CONTROL

simple and accessible

>> Find out how LabVIEW can make you better at **ni.com/labview/better**   **800 453 6202**

NATIONAL INSTRUMENTS

pulse
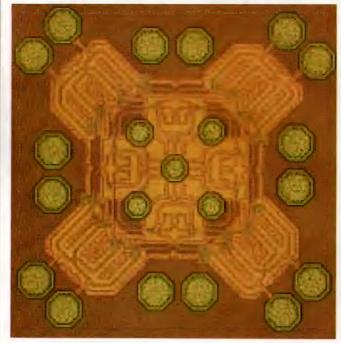
Terahertz scanner on your smartphone?

Terahertz radiation—the portion of the electromagnetic spectrum between microwaves and infrared light—can penetrate fabric and just a few millimeters into the skin without the damaging effects of X-rays, promising safer gear for medical and security imaging. To date, however, generating terahertz radiation has required lasers, vacuum tubes, and special circuits cooled near absolute zero, often in room-sized apparatuses costing thousands of dollars. An experimental method for generating terahertz signals using familiar, inexpensive CMOS chip technology could one day put the same functionality in a handheld device.

narrow frequency band. Connect two springs and set one vibrating, and the other also will begin to vibrate; eventually they will settle to equilibrium. A ring of electronic oscillators does the same, and the circuits coupling the oscillators can set the frequency at which they will lock in.

In Afshari's device, the couplers also shift the phase of the signals—that is, how the peaks and valleys of the waves line up. With the right adjustment, the peaks and valleys cancel out each other at several harmonics but reinforce each other at the fourth harmonic, channeling most of the power there.

The Cornell team report-

This electron-microscope image shows the prototype chip, which uses a ring of coupled oscillators to generate terahertz radiation. The design focuses most of the energy in a high harmonic; the signal radiates on the axis of the ring and can be aimed (courtesy Ehsan Afshari).

> **Cornell researchers fabricated experimental chips that generated signals with about 10,000 times the power level previously obtained at terahertz frequencies on a silicon device.**

Ehsan Afshari, assistant professor of electrical and computer engineering at Cornell University, came up with a tuning method that couples several oscillators in a ring to produce a high-quality signal, funneling the power into a very

edly fabricated experimental chips that generated signals with approximately 10,000 times the power level previously obtained at terahertz frequencies on a silicon chip. The Cornell researchers believe that by adjusting the couplers sepa-

rately, they might aim the output, making it possible to scan large areas with a narrow, high-powered beam.

A range of applications might harness terahertz radiation if the equipment cost could be contained. Terahertz scanners can identify skin cancers too small to see with the naked eye. Many of the complex organic chemicals used in explosives absorb terahertz radiation at

particular frequencies, creating a "signature" that detectors can read. And because higher frequencies can carry more bandwidth, terahertz signals might yield a "super Bluetooth" mechanism for transferring a full-length high-definition movie wirelessly in seconds.

Afshari is working with Cornell colleagues to implement his method in gallium nitride, which can handle both higher frequencies and higher power. The focus, however, is on CMOS chips that could deliver sophisticated scanning in a smartphone form factor.

—by Steve Taranovich
Cornell University,
www.cornell.edu.

DILBERT By Scott Adams

[www.edn.com]

pulse

Nordic offers BLE, ANT single-chip solutions

Nordic Semiconductor has announced the first members of its nRF51 SOC family for 2.4-GHz applications. Based on the ARM Cortex-M0 core, the devices include the multiprotocol nRF51822, targeting proprietary and BLE (Bluetooth Low Energy) applications, and the nRF51422, which Nordic says is the industry's first SOC for the ANT/ANT+ wireless network protocol. An improved radio design and the high-performance 32-bit core at the heart of each device enabled Nordic to boost the radio link budget by 9.5 dB while cutting power requirements.

Designed for low-power coin-cell battery applications, the devices feature peak current of less than 10 mA and a fine-grained power-manage-ment scheme. A programmable peripheral-interconnect design enables autonomous operation of peripherals, reducing the power consumption associated with processor operations; in fact, Nordic heavily leveraged the speed and increased code density of the 32-bit ARM core to reduce power-hungry CPU cycles.

"The Cortex's 100× faster start-up time means less current is required for start-up, and 10× higher performance means it finishes processing faster, so the duty cycle is much improved," says Thomas Embla Bonnerud, director of product management at Nordic. "As a result, the devices require 50% lower average current compared with previous generations of Nordic RF ICs."

In providing BLE and ANT/ANT+ stacks on the separate devices, Nordic used a protected memory architecture that separates application code from stack code. Developers use APIs provided as header files in the company's software development kit. All application code interacts with the provided stacks through function calls, a methodology that Nordic expects will speed development and reduce the conflicts often found in development of wireless applications based on software stacks provided as libraries.

According to Bonnerud, this separation was essential for reliable application development on SOC devices. "Two-chip solutions are good because you don't have to worry about the stack, but when you have to manage stack code in your application, you end up with tons of dependencies," he says.

In providing BLE and ANT/ANT+ stacks on the separate devices, Nordic used a protected memory architecture that separates application code from stack code.

"For the nRF51, the application developer sees a standard stack; but because we use clean APIs, the code is run-time-protected, so the application code can't do anything that affects the stack."

Sampling now to customers, the devices will be generally available in early September, with production quantities scheduled for availability later in the year. According to Bonnerud, pricing will be approximately $2, depending on quantity.

—by Stephen Evanczuk
▷**Nordic Semiconductor**, www.nordicsemi.com.

Printable conductive gel promises low-cost biosensors, energy storage

Stanford University researchers have reported on an electrically conductive hydrogel, synthesized from readily available organic material, that can be ink-jet printed and that demonstrates what researcher Yi Cui calls "remarkable electronic properties."

Polymer gels that feel and behave like biological tissue but are as electrically conductive as a metal or semiconductor hold promise for biosensor and energy-storage devices, but most of the gels developed to date are difficult to manufacture, and their conductive properties are limited. The Stanford team says its substance is inexpensive to make, and Cui states the gel's conductance is "among the best you can get" through the processes devised thus far.

Zhenan Bao, a chemical engineering associate professor, and Cui, a materials science and engineering associate professor, made the gel by binding long chains of the organic compound aniline together with phytic acid. The acid, found naturally in plant tissue, can grab up to six polymer chains simultaneously to yield an extensively cross-linked network that forms a complex, sponge-like structure. Pores in the nano-structured hydrogel expand its surface area, increasing the amount of charge it can hold, the speed of its electrical response, and its sensitivity as a chemical sensor.

Stanford postdoctoral fellow Guihua Yu, associate professor Zhenan Bao, and visiting scholar Lijia Pan examine the hydrogel synthesized in Bao's lab (courtesy Linda A. Cicero, Stanford News Service).

"There are already commercially available conducting polymers, but they all form a uniform film, without any nanostructures," says Bao.

The material doesn't solidify until the last synthesis step, so it can be printed or sprayed as a liquid that gels once it's in place, opening the door to low-cost manufacture of intricately patterned electrodes. "You can't print Jell-O," says Cui. "But with this technique, we can print it and make it Jell-O later."
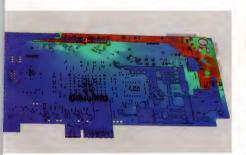
Most hydrogels are tied together by a large number of insulating molecules, reducing their ability to pass electrical current, but because phytic acid is a small-molecule dopant, the Stanford-developed gel is highly conductive. The substance's similarity to biological tissue, large surface area, and electrical capabilities suit it for application in biological systems that communicate with hardware, according to the researchers.

Stanford's Precourt Institute for Energy funded the research.

—by Diana Scheben
▷**Stanford University**, www.stanford.edu.

pulse

VOICES

Finally, a hardware accelerator: Bolt's Ben Einstein on a leg up for hardware start-ups

Experienced product designer and entrepreneur Ben Einstein is part of a team in Boston looking to give hardware start-ups their fair share of the glory. Bolt (http://buildatbolt.com), a new accelerator program that Einstein cofounded, aims to be a tool kit for hardware start-ups rather than for software and Web-services companies—popular draws for talent and investment that in the current environment have more support at their disposal. Einstein recently spoke to *EDN* about what Bolt is, what types of young companies it is looking for, and why the focus is on hardware. Read on for excerpts of the conversation, and get more of Einstein's thoughts at www.edn.com/4392154.

Tell us what Bolt is all about.
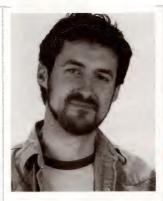
A We focus on the start-up community. We're basically an accelerator program designed to help hardware companies. The general premise is that if you are building a software company, there are so many resources available that are incredibly useful—software review, mentoring programs—and raising money is often fairly easy because there is so much awareness and popularity in the space. Bolt is explicitly built from the ground up to do similar things, but for hardware companies.

We have prototyping facilities and a full-time engineering staff; we help people go to Asia for manufacturing and tooling; we know how to talk to buyers. We have an open application process; anyone from anywhere in the world can apply. We will accept between 10 and 12 teams every six months and give them a little seed capital to help keep them alive.

Why focus on physical products as opposed to software or an app?

A There are maybe 150 [support programs] for software in the US [alone]. And there's this device phenomenon that is about to explode. [The targeted] companies act like software companies in terms of their revenue structure, the kinds of founders they attract, the way the IP is locked up, but they tend to have a piece of hardware that is sort of a barrier to entry.

There's a phrase that is thrown around, "software wrapped in plastic," which is what some of these companies begin to look like. They have a relatively simple piece of hardware, but it ends up enabling a new set of functionality and user applications [so that they] work like software companies; they have reoccurring revenue streams, multiple products around one piece of hardware, and reoccurring interactions with customers.

There's also [the fact that] these hardware companies aren't being helped at all. It's very difficult for some of these young companies to get the tools and systems they need to get to market, whereas in the app world a lot of that is pretty easy to come by. We are trying to leverage that.

What are you looking for?

A Great teams, but we are really looking for great people. The [online] application is tailored to find teams of great people. It's really simple; it's not designed to be a business plan.

When it comes to prototyping, what can an entrepreneurial engineer learn from a company like Dyson or iRobot or other former start-ups that showed distinct, difficult, or maybe risky designs in their beginnings?

A Bolt is based on building markets and helping people find their way along the path of building a product. It's less about the sexy, sleek prototype that people are really attracted to. The MVP— minimum viable product— approach involves doing as little as possible to prove that somebody really likes what you are doing. It's part of the lean start-up movement ... and involves growing the company around a core group of users and satisfying that need.

We take a similar approach when building products.

When people think about start-ups, they think about twenty-somethings just out of school. Do you see opportunities for more experienced engineers?

A We expect a good portion of our applicants to be in that [younger] category. But the other category is the frustrated engineers who are at a company like Apple or Motorola and have sort of been in their middle position and chugging along but feel that they want to be more creative and try out their own concept. Those [engineers] are *really* attractive to us. The only trick is that we are not providing enough capital to provide for families when the business is getting started. They have to be able to basically survive on their own for six months or so.

We expect that a smaller pool of the applicants would be these kinds of guys, but we encourage them to apply. You have to have a certain kind of crazy to start a company, especially in hardware, and that tends to be beaten out of people as they get older. But their experience [can be] incredibly valuable. Even if they are not leading the charge, it's great to have them on a team with two or three younger guys.—**interview conducted and edited by Suzanne Deffree**

designeast
center of the engineering universe

Einstein will be presenting at next month's Design East conference, hosted by UBM. Register and get more information at http://east. ubmdesign.com.

Protecting IGBTs with Avago Optical Isolation Amplifiers

Introduction

Insulated-gate bipolar transistors (IGBTs) can fail when subjected to overloads and overvoltages. Isolation amplfiers (iso-amps) can respond quickly to over-current and overload conditions when used on the output phases and the DC bus.

A typical block diagram of a power converter in an AC motor drivet consists of an inverter that converts the DC bus voltage to AC power at a variable frequency to drive the motor. IGBTs are expensive power switches that form the heart of the inverter. These power devices must operate at a high frequency and must be able to withstand high voltages.

Iso-amps such as the ACPL-C79A work with shunt resistors to accurately measure power converter current even in the presence of high switching noise. When used with a resistive divider, iso-amps work as precision voltage sensors to monitor the DC bus voltage. The microcontroller monitors the current and voltage information from the iso-amps and uses the data to calculate the feedback values and output signals needed to for fault management in the IGBTs and power converters.

Fault Protection

However, the IGBT protection must be such that its cost doesn't affect that of the motor drive system. IGBT gate drivers such as the ACPL-332J and current sensors with protection features can detect faults economically in this regard. They eliminate the need for separate detection and feedback components.

Over-current conditions in an IGBT can arise from a phase-to-phase short, a ground short or a shoot through. The shunt + iso-amp devices on the output phases and DC bus can, besides measuring current, detect such faults.
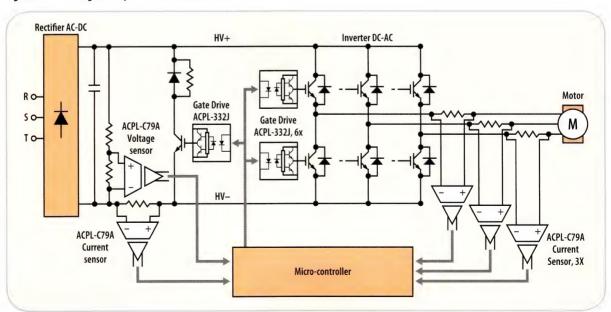
Typical IGBT short-circuit survival times are rated up to 10 μsec. So any protection must prevent this limit from being exceeded. Within 10 μsec, the circuit must detect the fault, notify the controller and complete the shutdown. Iso-amps use various methods to get these results.

For instance, the ACPL-C79A has a fast, 1.6 μsec response for a step input. That lets the iso-amp capture transients during short-circuits and overloads. The signal propagation delay from input to output at mid point is only 2 μsec, while it takes just 2.6 μsec for the output signal to catch up with input, reaching 90% of the final levels.

Another example is the HCPL-788J, which responds quickly to over-currents using a different approach. In addition to the signal output pin, it has a Fault pin that toggles quickly from High to Low level when over-current occurs. This iso-amp provides ±3% measurement accuracy.

In the fault feedback design, nuisance tripping can be an issue. This is a triggering of fault detection in the absence of any damaging fault condition. To avoid false triggering, the HCPL-788J employs a pulse discriminator that blanks out di/dt and dv/dt glitches. The advantage of this method is that

Figure 1: Block diagram of power converter in a motor drive
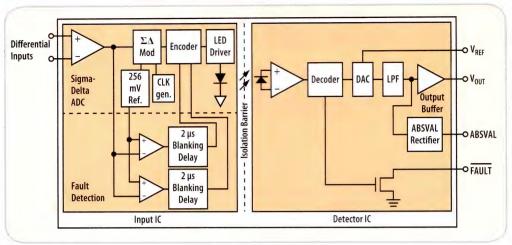
Figure 2:
In the HCPL-788J iso-amp, the differential input voltage is digitally encoded by a sigma-delta A/D converter and then fed to the LED driver, which sends the data across the isolation barrier to a detector and D/A.

rejection is independent of amplitude, so the fault threshold can be set to low level without risking nuisance tripping.

The circuit that detects faults quickly contains two comparators in the Fault Detection block to detect the negative and positive fault thresholds. The switching threshold is equal to the sigma-delta modulator reference of 256 mV. The outputs of these comparators connect to blanking filters with a blanking period of 2 μsec and then go to the Encoder block.

To ensure speedy transmission of the fault status across the isolation boundary, two unique digital coding sequences represent the fault condition, one code for negative, the other for affirmative. Detection of a fault interrupts the normal data transfer through the optical channel and replaces the bit stream with the fault code. These two fault codes deviate significantly from the normal coding scheme, so the decoder on the detector side immediately recognizes the codes as a fault conditions.

The decoder needs about 1 μsec to detect and communicate the fault condition across the isolation boundary. The anti-aliasing filter adds a 400 nsec delay to give a propagation delay of 1.4 μsec. The delay between the fault event and the output fault signal is the sum of the propagation delay and the blanking period (2 μsec) for an overall 3.4 μsec fault detection time.

The Fault output pin allows fault signals from several devices to be wire-ORed together forming a single fault signal. This signal may then be used to directly disable the PWM inputs through the controller.

Overload Detection

An overload condition refers to a situation where the motor current exceeds the rated drive current, but without imminent danger of failure, as when the motor is mechanically overloaded or is stalling because of a bearing failure

Inverters usually have an overload rating. The time period of the allowable overload rating depends on the time it takes before overheating becomes an issue. A typical overload rating is 150% of nominal load for up to one minute.

The ACPL-C79A accepts full-scale input range of ±300 mV and the data sheet specifications are based on ±200 mV nominal input range. Designers can choose the overload threshold at or in between either of the two figures. Usually the measurement accuracy of the overload current is less stringent than that of the normal operating current. Here, setting the threshold near 300 mV is a good choice. This allows full use of the iso-amp's dynamic input range. However, a threshold set at 200 mV ensures accurate measurement of the overload current. Once the voltage levels are decided, the designer must choose appropriate sense-resistor value according to corresponding current level.

The HCPL-788J includes an additional feature, the ABSVAL output, which can be used to simplify the overload detection circuit. The ABSVAL circuit rectifies the output signal, providing an output proportional to the absolute level of the input signal. This output is also wire OR-able. When three sinusoidal motor phases are combined, the rectified output (ABSVAL) is essentially a DC signal representing the RMS motor current. This DC signal and a threshold comparator can indicate motor overloads before they can damage to the motor or drive.

Overvoltage Detection

The DC bus voltage must also be continuously controled. Under certain operating conditions, a motor can act as a generator, delivering a high voltage back into the DC bus through the inverter power devices and/or recovery diodes. This high voltage adds to the DC bus voltage and puts a very high surge on the IGBTs. That surge may exceed the maximum IGBT collect-emitter voltage and damage them.

The miniature iso-amp (ACPL-C79A) is often used as a voltage sensor in DC bus monitoring applications. A designer must scale down the DC bus voltage to fit the input range of the iso-amp by choosing R1 and R2 values to get an appropriate ratio.

AVAGO
T E C H N O L O G I E S

BY BONNIE BAKER

Instrumentation amplifier input-circuit strategies

Many industrial and medical applications use instrumentation amplifiers to condition small signals in the presence of large common-mode voltages. INAs have complete closed-loop operational amplifiers with feedback components included. Under normal conditions, an INA is easy to use, as long as you pay attention to the input stage and the output of the first stage. The quick and inexpensive solutions described here can help you avoid the INA circuit's input-stage pitfalls.

In a classic, three-op-amp INA (**Figure 1**), the input stage has two op amps in an adjustable-gain configuration and provides high input impedance on both the inverting (V_{IN-}) and noninverting (V_{IN+}) inputs. The output stage has four matched resistors around a single op amp. When the circuit designer exercises proper precautions with the input pins, this configuration rejects external common-mode voltage and noise. The output stage has a reference voltage that level-shifts the output with respect to ground. The level shift is convenient in single-supply applications. As with many INAs, you can program the gain found in **Figure 1** with a single resistor, R_G.

Depending on the INA's silicon process, V_{IN+} and V_{IN-} connect to a bipolar transistor base, FET gate, or CMOS gate. All inputs to the INA require a current-return path to ground and a bias-voltage reference. Without the current-return path and the bias-voltage reference, the INA input stage saturates or floats to an undesirable voltage. Either condition creates an invalid output voltage.

The floating thermocouple circuit in **Figure 2a**, available online at www.edn.com/4392026, does not provide a current path to ground or a bias-voltage reference for the INA's input pins. Thus, the input-current leakages are not dissipated, and the two inputs can float to any undefined voltage. That situation, in turn, causes the INA's output to change to an invalid output voltage, usually in the middle of the INA's output range. The invalid voltage can appear to be legitimate, making it difficult to detect the correctness of the INA input implementation. **Figure 2b** illustrates a correct thermocouple connection to an INA; both inputs have a path to ground—in this case, through a 10-kΩ resistor, biased to a voltage within the INA's input range, or ground.

Circuit designers often misapply a thermocouple or even a two-wire microphone to the INA's input circuit; the problem arises when the INA inputs are connected without proper consideration to current paths or biasing. Heed those considerations for the INA's input stage, and you can be confident that the INA's output voltage is representative of the thermocouple's voltage.

But wait—there may be a problem with the output values of A_1 and A_2 (**Figure 1**). How would *you* solve this problem? Comment online at www.edn.com/4392026. I'd really like to hear from you! **EDN**

Figure 1 In a classic, three-op-amp INA topology with input sources that define V_{CM} and V_D, the input stage has two op amps in an adjustable-gain configuration and provides high input impedance on both V_{IN-} and V_{IN+}.

$$G = 1 + \left(\frac{100 \text{ k}\Omega}{R_G} \right)$$

REFERENCES

[1] Kugelstadt, Thomas, "Getting the most out of your instrumentation amplifier design," *Analog Applications Journal*, Texas Instruments, 4Q 2005, http://bit.ly/Peilok.

[2] Trump, Bruce, "Instrumentation amplifiers—avoiding a common pitfall," TI E2E Community, March 27, 2012, http://bit.ly/Nw4er2.

[3] Trump, Bruce, "Instrumentation amplifiers—they're not op amps but what are they?" TI E2E Community, April 2, 2012, http://bit.ly/NCsxHI.

Bonnie Baker is a senior applications engineer at Texas Instruments.

Volt's battery-stack manager aids Chevy's drive for safe power
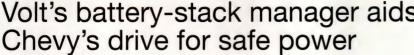
A sophisticated battery-management system at the heart of the Chevrolet Volt ensures the safety and reliability of the multicell lithium-ion battery stack that delivers power on demand to the Volt drive system. Within the management system, battery-monitoring boards use two key subsystems to monitor cell health and deliver digital results to a host processor that orchestrates system operation. Separating those subsystems, a signal interface ensures isolation between high-voltage battery-sensing circuitry and communications devices on the boards.

At a given temperature and output current level, Li-ion cells maintain a nearly flat voltage output across the middle of their capacity range—a characteristic that heightens Li-ion's advantages as a power source but complicates engineers' attempts to use simple cell-voltage measurements to provide users with a measure of remaining battery capacity, or SOC (state of charge).

For Volt drivers, an accurate SOC measurement is important for estimating vehicle range. Maintaining SOC within a specific range is also important for ensuring extended battery life; batteries kept at too low or too high a state of charge tend to degrade faster than those kept at intermediate levels.

In the Volt, GM engineers established a safe SOC window of 58% to 65%, depending on driving mode, with a lower limit in normal driving mode set to a 30% SOC and a higher, "mountain driving" mode limit set to 45%. When the Volt reaches the appropriate lower SOC limit, the vehicle's gasoline engine engages for extended-range operation. The battery-interface control-module boards provide the data needed to ensure a reliable SOC for Volt operation.

Measuring 1.8m long and weighing 181 kg, the battery stack generates the 16-kWh power needed to turn drive motors, power passenger features, and supply power to a battery-management system that rivals avionic systems in its complexity. As described by analysts at UBM TechInsights and Munro & Associates, the full Chevy Volt battery pack comprises 288 prismatic Li-ion cells, which are in turn packaged into 96 battery-cell groups to deliver a system voltage that the analysts measured at 386.6V dc.

The battery-cell groups are combined with temperature sensors and cooling elements into four main battery modules. Voltage-sense lines attached to each cell group terminate in a connector on top of each battery module, where a voltage-sense harness joins the connector to a battery-interface control module.

The orange battery connectors carry the battery-module temperature-signal, low-reference, and voltage-sense lines from the battery-cell groups.

The battery-interface control module is a four-layer PCB with most of the components mounted on the top, along with orange battery connect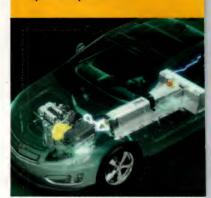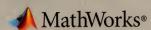ors and a black data-communications connector. The uppermost layer includes a ground plane and signal traces with multiple vias providing connections to lower layers. In the second layer, power and ground planes spread under the high-voltage areas of the PCB; the third layer contains signal traces passing under those areas. The other side of the PCB, the fourth layer, is used for the ground plane and signal traces and contains some additional components.
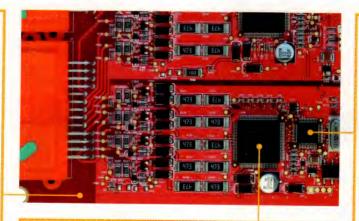
The black ATLPB-21-2AK PCB-mounted connector carries 5V reference, low-reference, signal-ground, CAN bus high-serial-data, CAN bus low-serial-data, and high-voltage-fault signals.

The heart of the battery-interface control system is a sophisticated sensing subsystem—a complete embedded system for monitoring the battery-pack temperature and the voltage output from each Li-ion cell group. Each sensing subsystem pairs an L9763 ASIC, developed by STMicroelectronics and LG Chem, with a Freescale S9S08DZ32 40-MHz HCS08 microcontroller packing internal 32-kbyte flash, 2-kbyte RAM, and 1-kbyte E2PROM. An external, 4-MHz oscillator provides a reference frequency for microcontroller clock operations.

The S9S08DZ32 receives measurement data from the L9763 ASIC, which passes measurement data for its monitored Li-ion cells through its SPI interface to the microcontroller. The ASIC also provides a 5V LDO output for the microcontroller. For overall battery management, individual L9763 ASICs are linked through on-chip interfaces, and the primary control unit addresses them individually through vertical daisy-chain communication.

Cell voltage passes through the battery connector to the L9763. The ASIC can monitor up to 10 individual Li-ion cell groups, using on-chip current-sense amplifiers for cell-load-current monitoring and an on-chip analog multiplexer and sample-and-hold circuit for cell-voltage monitoring. The part's differential inputs address the need to ensure millivolts-accurate measurements despite large offset voltages, depending on the cell's location in the battery pack. PCB designers used trace-layout techniques, isolation techniques, and the ground planes mentioned earlier to ensure signal integrity.

In the battery-interface-module PCB, Avago ACPL-M43T optocouplers sit on the edge of the communications section, isolating it from the high-voltage-sensing subsystems, which are further shielded by ground planes in deeper PCB layers. The isolation interface provides three individual M43T optocouplers for each of the three lines brought out from each sensing circuit—specifically, the Freescale S9S08DZ32 CAN Tx output pin, microcontroller CAN Rx input pin, and high-voltage fault signal from the microcontroller. Output from the microcontroller CAN Tx pin, for example, passes through a shielded signal layer in the PCB to reach the pin 1 anode of an M43T to energize the embedded LED, resulting in a change of state at pin 5 (V_o). The isolated signal then passes to the communications output stage of the battery-interface module.

Sitting at the end of the communications signal chain, the Infineon TLE6250G CAN transceiver is an AEC-qualified IC that provides CAN physical-layer signaling between the physical cable and CAN protocol handler—here, the S9S08DZ32 microcontroller (via optocoupler isolation). Rated for 1M-baud CAN transmission, the device handles the conversion between the CAN_H and CAN_L signals on the differential signal wire and the CAN dominant and recessive bits transmitted and received by the S9S08DZ32. The eight-pin TLE6250G includes Tx, Rx, V_{CC}, GND, CAN_H, and CAN_L pins, as well as two mode-control pins: INH and RM. When the TLE6250G senses a change on the Rx pin from the CAN recessive state to CAN dominant state, the device switches CAN_H high and CAN_L low. This symmetrical change in state reduces electromagnetic interference.

BY PALLAB CHATTERJEE

In wireless, MIMO's the word

The fastest growing avenue of communications is via the wireless technologies that drive cellular communications, Wi-Fi, wireless display interfaces, and the large number of machines in the Internet of Things. This explosive growth is not without technical challenges, the most stringent of which is spectrum allocation. The more available channels and frequencies there are, the higher the number of devices that can communicate at one time, and the greater the bandwidth available for each device.

One driver for increased bandwidth use is video traffic over Internet Protocol networks. Video currently claims 60% to 75% of the IP bandwidth for Internet traffic overall, and it accounts for upward of 80% of the traffic on home networks and in industrial-control applications (including surveillance- and vision-system traffic). This video data includes still images as well as streaming and multiframe content.

One change to support the growing volume of traffic has been an increase in the frequency bands available for wireless connections. Wireless spectrum initially resided in the 900-MHz band, expanded to the 1.2-GHz and 2.4-GHz bands, and has now landed in the 5-GHz band.

On the mainstream Wi-Fi side (802.11), the a/b/g protocols are situated in the 2.4-GHz band and feature single antennas with single data streams. The 802.11n protocol supports dual-band, 2.4-GHz and 5-GHz operation; more important, it implements MIMO (multiple-input/multiple-output) data streams. Not only does the higher frequency help, but allowing two streams via the dual-transceiver configuration eases the transfer of large data sets.

Figure 1 shows the progression in antenna techniques that have been used. The MIMO configuration for 802.11n was primarily seen in routers and gateways, rather than endpoints, in the first round of devices.

The concept was expanded and optimized for the new 802.11ac protocol, which allows for a four-channel MIMO configuration at a 5-GHz fixed frequency. Most ac-class routers will support the multiple-antenna set. A key change is that endpoint devices—tablets, smart-phones, laptops, desktop devices, printers, and displays—can now support multiple antennas and multiple data streams. With the increased bandwidth from MIMO in the endpoint and router, high-definition and high-frame-rate video can now be transferred between devices without any QOE (quality of experience) lag or degradation.

Even as 802.11 Wi-Fi has shifted to a MIMO platform for new products, machine-to-machine protocols such as 802.15.4 under ZigBee and Z-Wave have done the same. The M2M protocols support dual configurations for the points in a mesh network. The dual transceivers typically function not as a pair to increase the effective transfer rate but as a simultaneous-send/receive pair. That setup enables video-data transfer through the mesh network to the final endpoint, while small data is transferred between points from the sensors that originated it.

Moving to a MIMO configuration involves other system changes for improved operation. For wireless systems, developers must coordinate the data and signal flows through the multiple antennas in the product to improve performance.

The technique of beamforming can improve data transfer in the multiple-broadcast configuration. Beamforming adjusts the output-signal phase to compensate for the phase shift that occurs during the broadcast transmission. The technique ensures a high total signal level because the combined output signals do not cancel one another as a result of phase overlap.

Another major issue for a MIMO implementation is power. A setup that allows multiple antennas to be active at the same time requires separate RF blocks connected to the antennas. A four-channel system would require four power amplifiers and transceivers, with each capable of delivering the full power of the protocol set.**EDN**

Figure 1 Wireless-antenna configurations have progressed from single-input/single-output to multichannel MIMO.

Pallab Chatterjee has been an independent design consultant since 1985.

ELECTRON-SIPPING

COUNTS
ON
CONNECTIVITY

THAT'S WHY THE WORLD'S AUTOMOTIVE COMPANIES COUNT ON TE CONNECTIVITY.

Soon, the ultimate connected device won't be in your pocket. It will be on the road. And TE Connectivity is making it possible, with connectivity solutions that drive today's most advanced hybrid and electric vehicles. We're helping automotive engineers develop safer, lighter, more robust ways to link power and data—from the Smart Grid to the charging station and throughout the vehicle. When the world needs a better way to get where it's going, we make every connection count. **EveryConnectionCounts.com**

TE
connectivity

EVERY CONNECTION COUNTS

MECHATRONICS
IN DESIGN

FRESH IDEAS ON INTEGRATING MECHANICAL SYSTEMS, ELECTRONICS, CONTROL SYSTEMS, AND SOFTWARE IN DESIGN

"Mechatronify" common mechanisms

Through a combination of knowledge, process, and determination, even a Renaissance-era mechanism can be made mechatronic.

By Kevin C Craig, PhD

What does it mean to "mechatronify" a mechanism or machine? It's more than just adding a sensor, actuator, and computer-control system; you must add them in an integrated way from the start of the design, using a model-based process that leverages analysis techniques and simulation software and that leads to optimal designs without trial and error. Combining old inventions with new technology fosters innovative ideas, but it is the process that transforms those ideas into reality.

Consider a mechanism that Leonardo da Vinci developed more than 500 years ago and that now finds use in engines, automation applications, and miniature devices: the slider crank (**Figure 1**). The slider-crank mechanism comprises a flywheel crank, a connecting rod, and a slider, all assumed to be rigid. The external forces and torques acting on the mechanism are the servo-motor torque, τ; the slider friction force, F_f; and the external force, F_e. In this special case of four-bar linkage, one crank is infinite in length, such that its end point (B) has rectilinear motion. It is a one-degree-of-freedom system; that is, only one coordinate is needed to describe its motion completely. The constraint **equation** relating angles θ and φ is $r(\sin \theta) = \ell(\sin \varphi)$.

Graphical analysis of kinematics—the geometry of motion—can yield great insight. A mathematical solution, however, is more effective for mechatronic system design and optimization, obtaining positions, velocities, and accelerations of key points, as well as the angular velocities and accelerations of the rigid bodies.

You can accomplish kinetic analysis by drawing free-body diagrams showing gravitational forces, contact forces/torques, and inertia forces/torques and then summing forces/moments as needed. This method, known as the

D'Alembert approach to applying Newton's Laws, can determine forces/torques at all joints. You can directly obtain the system **equation** of motion by applying Lagrange's Equation:

$$\frac{d}{dt}\frac{\partial T}{\partial \dot{\theta}} - \frac{\partial T}{\partial \theta} + \frac{\partial V}{\partial \theta} = Q_\theta,$$

using the generalized coordinate θ, where V is the system potential energy, T is the system kinetic energy, and Q_θ is the generalized torque due to forces/torques that do work. The resulting **equation** has the form $M(\theta)\ddot{\theta} + N(\theta, \dot{\theta}) = F(\theta)$.

Having completed the kinematic and kinetic analyses, you would next define the desired end-point trajectory and then compute the necessary motion profile for the actuator through inverse kinematics, which here would include the crank and connecting rod lengths. This step involves trajectory planning; you must define the profile in a way that avoids or reduces the mechanical vibration and stress on components and actuators while reducing overshoot response and excessive position error during motion. Electronic cams accomplish this task.

The inverse kinetic analysis, which includes masses, center-of-mass locations, and mass moments of inertia, is used to generate the required actuator torque/force for the motion profile, and results in a speed/torque-force diagram on which to base actuator selection. The chosen actuator now becomes part of the system; with that update, a control system—including feedback, feedforward, and disturbance observer—can be designed, yielding a new speed/torque-force profile. You would then simulate the entire system, with the addition of any parasitic effects, for design validation.

Mechatronifying any mechanism requires sweating the details. But when you apply a combination of accumulated knowledge, process, and determination, innovation will happen.**EDN**
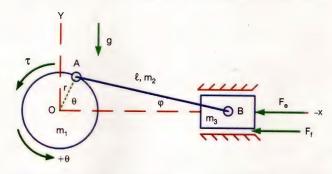
Figure 1 The slider-crank mechanism—a concept dating to da Vinci—comprises a flywheel crank, a connecting rod, and a slider.

MIXED-SIGNAL
SOC VERIFICATION
USING ANALOG BEHAVIORAL MODELS

BRINGING ANALOG AND MIXED-SIGNAL BLOCKS TO A HIGHER ABSTRACTION LEVEL YIELDS MORE EFFECTIVE MIXED-SIGNAL SIMULATION AND MORE COMPLETE VERIFICATION ENVIRONMENTS.

BY QI WANG · CADENCE DESIGN SYSTEMS

The era of "Internet everywhere" is creating a spectrum of applications targeted toward low-power and mixed-signal design, in segments ranging from health care to automotive to communications. Meanwhile, design challenges such as intellectual-property selection and integration as well as SOC- and system-level verification are spawning a whole new class of problems for EDA tools.

Mixed-signal design engineers face increasing difficulties in design and verification of complex mixed-signal SOCs. In a survey of mixed-signal design engineers during the 2011 Mixed-Signal Tech on Tour, a worldwide series presented by Cadence Design Systems Inc, the 561 respondents identified mixed-signal verification as a top customer challenge.

The performance of Spice simulation is prominent in the difficulties being reported (**Figure 1**). Analog Spice and Fast-Spice simulators are orders of magnitude slower than digital simulators and are slower still when compared with emulators and hardware accelerators. A June 2011 Design Automation Conference panel discussed the need for analog design and verification to become more like digital—that is, to become more structured and more top-down (**Reference 1**). Verification planning tools are required, and debug methodologies such as ABV (assertion-based verification), MDV (metric-driven verification), and UVM (universal verification methodology)-like self-checking test benches must be created for analog/mixed-signal.

To tackle simulation-throughput issues, designers are turning to behavioral-modeling techniques, which can increase simulation speed. Such techniques include event-driven simulation based on Verilog-A, Verilog-AMS, and RNM (real-number modeling).

Analog behavioral models are typically written in Verilog-AMS, Verilog-A, VHDL-AMS, or SystemVerilog.

Verilog-A is a pure-analog subset of Verilog-AMS and is mainly used for detailed analog models for performance verification. The language is quite simple, but it is challenging to write a good behavioral model with Verilog-A that provides significant performance

AT A GLANCE

☒ Mixed-signal design engineers face increasing difficulties in design and verification of complex mixed-signal SOCs.

☒ Real-number modeling defines blocks in terms of input/output transfer characteristics, with no strong direct feedback present among the blocks.

☒ The digital world has many tools at its disposal—such as verification planning, coverage, and assertions—that can apply to analog verification given the right language extensions.

gains while retaining the right level of accuracy. The advantage of Verilog-A is the ability to use models in pure-analog simulations as well as in the mixed-signal environment. The models are too low-level, however, to enable efficient SOC-level verification of mixed-signal designs.

The RNM technique models electrical signals by representing them as real values. Provided that the modules are at a sufficiently high level of abstraction, the interfaces can be described by passing real numbers between blocks to represent the voltage, or current, signal being transferred. This is a powerful way to simulate complex systems rapidly. Traditionally, iterating to a solution involving feedback would require

an analog solver (see **sidebar**, "Analog versus digital solvers").

RNM is available in the Verilog-AMS, SystemVerilog, and VHDL-AMS languages. A commonly used RNM approach is the wreal data type in Verilog-AMS. RNM uses a discrete event solver—without an analog solver—and can be used to simulate mixed-signal systems at incredible speeds. It is primarily limited to modeling at a high enough level of abstraction that bidirectional analog interactions between blocks are not significant. In other words, typical RNM defines blocks in terms of input/output transfer characteristics, with no strong direct feedback present among the blocks. Logic can be modeled naturally in these languages, so RNM is also a good choice for systems with only a small amount of analog content.

TOP-DOWN OR BOTTOM-UP

Designers use two principal methodologies based on the creation of behavioral models for mixed-signal design. In a top-down methodology, models are developed before the circuits are designed. The behavioral models can be simpler ones that are sufficient for functional verification at the system level. In a bottom-up methodology, the models are written to match an already implemented block for performance verification and usually result in a more accurate but slower-running model.
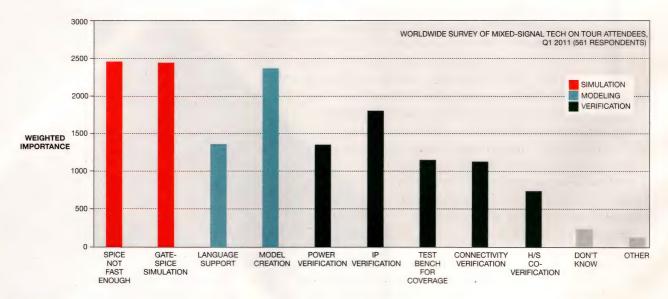
Figure 1 Results of a 2011 mixed-signal survey show that verification-related issues figure highly among customer challenges.

The creation of analog behavioral models can be challenging. Analog designers are in the best position to create such models because they are familiar with their own circuits. Many analog designers lack the programming skills and knowledge required to construct behavioral models, however, and few are familiar with Verilog or VHDL. Digital designers, conversely, have expertise with behavioral models but know less about analog circuits.

That dichotomy creates an opportunity for tool vendors to offer an automated or semi-automated solution for generating analog behavioral models. For example, automated, netlist-driven model-generation technologies can create a fairly accurate parametric behavioral model that considers PVT (process, voltage, and temperature) and loading variations for functional verification. Such an approach has shown some limited success on a subset of analog-circuit architectures under specific contexts, but there is still much work to be done to develop a general model-generation methodology with high accuracy that can be applied to any analog or mixed-signal design.

Creating behavioral models is only one part of the process of using those models in a mixed-signal verification flow. If the model and implementation do not match, the effort is worthless; worse, it can damage the entire design process. As a result, there is a need for a methodology to validate the accuracy of a behavioral model automatically against the corresponding design. The model also must be updated to keep it in sync with any changes made in the implementation.

The model-validation flow shown in **Figure 2** simulates the implementation and behavioral models using the same test bench, with the relevant tools and flows creating the required measured-results behavior and waveform for each model. The designer can use the flow to validate that:

• The implementation- and behavioral-model measured results are within user-provided tolerances;

• The implementation- and behavioral-model waveforms are within user-provided tolerances; and

• Selected elements of the interfaces and pins of the behavioral and implementation models match.

This approach provides a self-contained, automated analog-behavioral-model verification and debugging environment. The method lets users provide verification requirements through interactive configuration, and it uses the existing mixed-signal simulation setup to validate both the waveform signals and the measured results. Interactive verification-failure debug lets designers quickly identify problems in a model and rerun validation after fixing them.

That is just the beginning of the advantages that can be brought into the analog-verification world. The digital domain has many other tools at its disposal—such as verification planning, coverage, and assertions—that

can apply to analog verification given the right language extensions.

ASSERTIONS' ROLE

Assertions capture aspects of a specification and design intent in an executable form. They act as monitors during simulation, detecting errors close to their source and reporting both errors and coverage information. Through the use of assertions, verification can start earlier; design and verification teams can detect and remove bugs faster; and designers can incorporate their intent into the design code, thereby minimizing integration issues later on. The mixed-signal design and verification communities embracing ABV methodologies can reap the benefits of assertions and tools in the following ways:

• Assertions capture design intent and can be incorporated with the design;
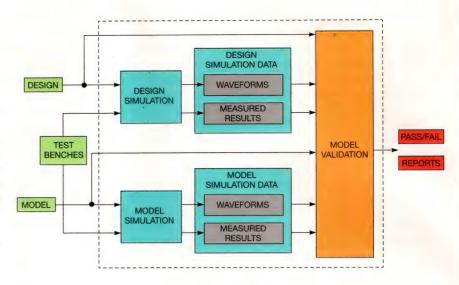
Figure 2 This mixed-signal-model-validation framework uses the same test bench to simulate the implementation and behavioral models.

ANALOG VERSUS DIGITAL SOLVERS

Digital simulators execute much faster than their analog counterparts. Digital simulators calculate values for outputs whenever inputs change. There is a time delay associated with those value changes. Each change is independent of other changes happening in the system.

Analog simulators are tasked with solving a set of simultaneous (nonlinear, ordinary differential) equations. During transient analysis, matrix-based numerical-analysis techniques solve the complete set of voltage and current equations at each analog time step. The process is iterative because it involves feedback; performance depends in part on how quickly the simulator can converge on a solution.

A direct-solution method could immediately compute the voltages and currents at each time point without requiring an iterative process if the method had to define only linear systems, but real systems are never completely linear. Everything has bounds of operation, often induced by power-supply or performance limits that result from the transistors used to implement the characteristics.

The typical analog simulator breaks the time axis into sufficiently small time steps to allow it to approximate linearity over each of the steps (Figure A). It then can calculate the simultaneous-equation solution that describes what should happen over each small time step. A significant amount of theory is associated with doing this calculation accurately, predicated by the simulator's ability to estimate how much error may occur at each step. The simulator can then decide how big a time step it can safely take, and it must iterate and converge toward a solution that sufficiently solves the Kirchhoff's voltage and current laws at the new time point.

With a nonlinear system, this iteration procedure is not guaranteed. Each iteration estimate of the new solution requires reevaluation of the nonlinear relationships in the system. If the solution does not converge to a sufficiently accurate value in a certain number of iterations, the simulator will typically assume that the step size was too large for the nonlinearities and give up computing that point. It will then try a smaller time step from the previous time point.
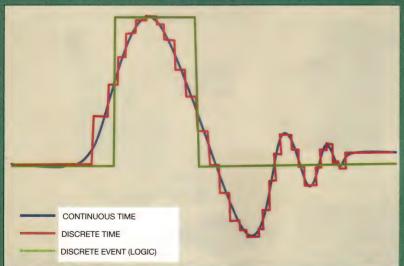
Figure A Various sampling methodologies are compared. Analog simulators break the time axis into small time steps.

- Assertions detect errors closer to their source, speeding removal; and
- Assertions provide control-oriented functional coverage information.

Assertions can also capture properties about interfaces, such as the assumptions made in the design, the interactions that are expected to occur (or not occur), and errors related to behaviors such as nominal functional boundary conditions or startup. A metric-driven verification environment can use assertions as coverage points, which ensure that the environment has verified the design in every valid configuration, that it has verified all possible communications-protocol variations, that valid combinations of inputs have been applied, and that valid combinations of outputs have been observed.

In this still-emerging field, some HDLs and HVLs include support for real-numbered functional-coverage objects; the ability to generate real-valued, constrained random numbers; and the ability to create assertions on electrical quantities, such as voltages or currents.

Two standards groups are actively working to standardize analog/mixed-signal assertions. The ASVA (Analog System Verilog Assertions) committee is targeting analog and mixed-signal extensions to the SVA subset of SystemVerilog (**Reference 2**); the SV-AMS (SystemVerilog-AMS) committee is defining analog/mixed-signal extensions to SystemVerilog in work that parallels the successful transformation of Verilog into Verilog-AMS (**Reference 3**). Participants expect the results of the ASVA committee's work to feed into the longer-term SV-AMS effort.

Behavioral modeling continues to make inroads into mixed-signal simulation and verification as a result of behavioral-language standardization, skills development, and automation improvements. As design teams gain a better understanding of the advantages that behavioral modeling offers, more are making the initial investment necessary for the methodology shift, confident they will reap a return in the form of productivity improvements.

Behavioral modeling can be a key component of any mixed-signal verification methodology. Bringing analog and mixed-signal blocks to a higher level of abstraction enables the creation of more effective mixed-signal simulation and more complete verification environments.EDN

REFERENCES

1 Goering, Richard, "DAC Panel: Users Describe Mixed-Signal Verification Challenges, Solutions," Cadence Design Systems Inc, June 13, 2011, http://bit.ly/PQxkax.

2 Verilog Analog/Mixed-Signal Technical Subcommittee, http://bit.ly/PMgo3B.

3 SystemVerilog-AMS, http://bit.ly/NsPMBZ.

AUTHOR'S BIOGRAPHY

Qi Wang is group director for solutions marketing at Cadence Design Systems, with a focus on low-power and mixed-signal solutions. In this role, he works with product-marketing, research and development, and technical-support teams from various product groups to set the solution-wide strategy and to promote Cadence solutions for advanced low-power and mixed-signal designs.
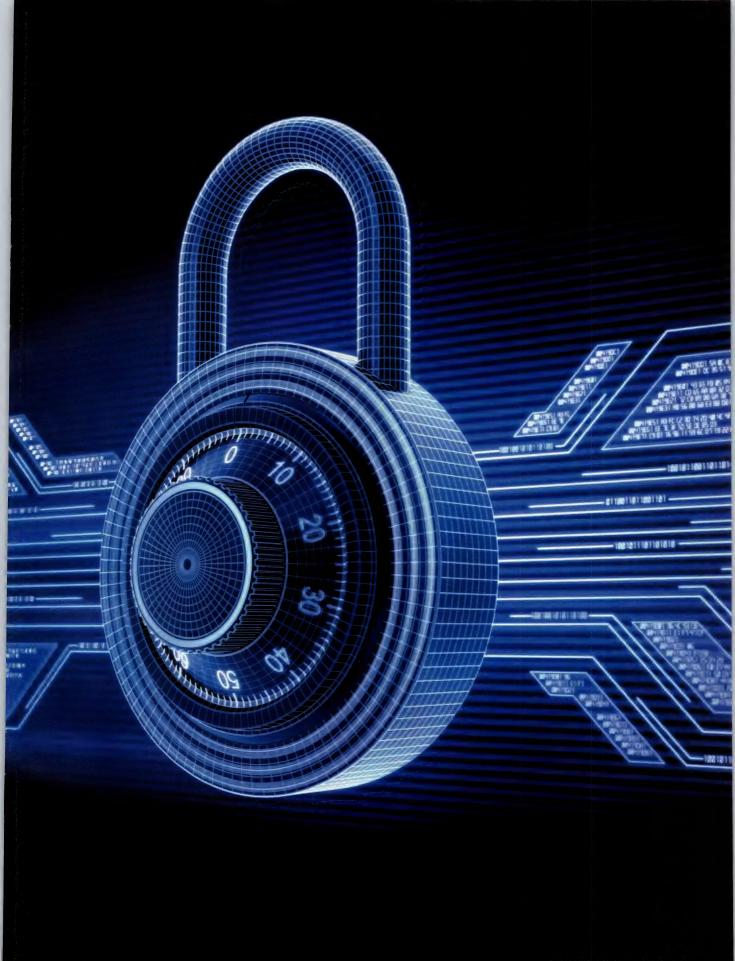
THE FUTURE OF
CONNECTED-DEVICE
SECURITY

DESIGNERS BUILDING SCALABLE DATA
PROTECTION INTO NEXT-GENERATION
ELECTRONICS CONFRONT A PLETHORA OF
DESIGN CONSTRAINTS AND TECHNOLOGY
CHOICES, BUILDING UP FROM THE HARDWARE
AND SOFTWARE ROOTS OF TRUST.

BY DAVID KLEIDERMACHER · GREEN HILLS SOFTWARE

The electronics world is seeing rapid growth in sophistication, driven by M2M intelligence, multimedia capability, Internet connectivity, and high-value financial transactions. Such capabilities imply juicy network attack vectors and are attractive targets for hackers, generating security requirements that electronics designers must learn and embrace. Here, we offer practical implementation guidance, focusing on hardware and software roots of trust, data-storage protection, and secure network connectivity.

IMAGE: ISTOCK

SECURITY TRENDS

One of the first electronic systems within an automobile was the 1978 Cadillac Seville's trip computer, run by a Motorola 6802 microprocessor with 128 bytes of RAM and 2 kbytes of ROM. The printed source code could not have occupied more than a handful of pages.

Today, even the lowest-end automobile contains at least a dozen microprocessors; luxury cars are estimated to contain approximately 100 microprocessors. With infotainment systems running sophisticated operating systems such as Windows and Linux, the total electronics software content can easily exceed 100 million lines of code.

This electronics evolution has been beneficial to society but has also served as a key source of security woes. Many problems relating to losses in quality, safety, and security in electronics are rooted in rising complexity that cannot be effectively managed.

Another clear trend in electronic systems is the addition of network connectivity, enabling, for example, remote management and the ability to upgrade software in the field. In 2010, General Motors introduced a feature to let car owners manipulate the locks and start the engine using a smartphone. Just before GM announced the smartphone feature, however, a team of university researchers published a study demonstrating how potential attackers could exploit vulnerabilities in such a car's electronics to tamper with its critical systems, such as brakes and engine throttling (**Reference 1**). Researchers now are demonstrating attack vectors—such as the telematics connection—over wide-area networks.

The increasing reliance on electronic systems in commerce, critical infrastructure, and life-critical function makes them attractive to well-funded and determined attackers. Industrial control systems managing nuclear reactors, oil refineries, and other critical infrastructure present opportunities to inflict widespread damage.

The Stuxnet worm, which infiltrated Siemens process-control systems at nuclear plants, was likely the first malware to target electronic process-control systems directly. Stuxnet demonstrates the incredible sophistication of modern electronics security attacks

AT A GLANCE

⤵ Increases in design complexity, connectivity, and the sophistication of attacks are shaping the future of connected-device security.

⤵ Security-conscious designs begin with both a hardware and a software root of trust: a tamper-proof mechanism to ensure trusted firmware is always loaded, as well as a trustworthy software base that can host security functions and protect both them and itself from the more sophisticated, noncritical portions of the system.

⤵ Embedded computer system virtualization shows promise for providing the software root of trust while enabling all the bells and whistles demanded of next-generation connected devices.

⤵ Building on the root of trust, designers must consider the proper approach for incorporating scalable data protection: secure connectivity for functions such as trusted field upgrade and remote management, as well as sensitive data-at-rest encryption to protect intellectual property and private information.

and illustrates the need for electronics designers to hone their security skills.

Another trend affecting security is processor consolidation. In the modern automobile, for example, skyrocketing electronics content poses production-cost, physical-footprint, and time-to-market challenges for manufacturers. The response is to reverse the trend and merge disparate functions into fewer electronic components. Consolidation requires the proper systems architecture to ensure that those components do not interact in unforeseen ways, posing a reliability risk.

For example, automotive-system designers are consolidating the infotainment head unit with the rear-view camera and ADAS (advanced driver-assistance systems). Because it can share the center stack computer's audio and video capability, the rear-view camera module is a natural candidate for consolidation, yet it is considered a safety-critical function (**Figure 1**).

These trends all point to a need for designer security training and the incorporation of critical security technologies to help manage complexity and provide protection.

ROOT-OF-TRUST CONCERNS

Electronic-system security should begin with a design that provides confidence to users that the known, tested firmware or software is always running and has not been tampered with during delivery or when fielded. The act of establishing a secure initial state is often referred to as secure boot.

In electronic systems, a chain of firmware must be executed to establish the secure initial state. Usually, the CPU first executes a small boot loader, burned into ROM at manufac-
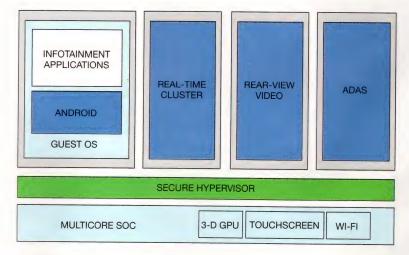

Figure 1 Next-generation automotive electronics must support high-end multimedia operating systems and safety-critical, real-time applications on the same computer.

turing time. Secure boot depends on a hardware-based root of trust; in this case, it is the fact that the ROM cannot be modified postproduction.

The ROM loader will often boot a more functional, second-level boot loader residing in internal flash. In many cases, the second-level boot loader will launch the primary operating system, which in turn starts the OS applications.

The typical secure-boot method is to verify the authenticity of each component in the chain. If any link is broken, the secure state is compromised. The first-stage loader must have a preburned cryptographic key that verifies the digital signature of the next-level boot loader. The key may be integrated into the ROM loader image itself, installed using a one-time-programmable fuse, or stored in a local TPM (trusted platform module) that may provide enhanced tamper protection. The hardware root of trust must include this initial verification key.

The signature key verifies the authenticity of the second-stage component in the boot chain; therefore, the hardware-protected area must also store the known-good signature. Verification of the second-level component covers its executable image as well as the known-good signature and signature-verification key of the third stage, if there is one.

The chain of verification can be indefinitely long; it is not uncommon for some sophisticated systems to have surprisingly long chains or even trees of verified components that make up the TCB (trusted computing base). In computer-security terms, TCB refers to those portions of a system, both software and hardware, that are critical to security and that therefore must be trustworthy. **Figure 2** depicts a three-level secure-boot sequence.

Secure boot provides electronics designers with confidence that the product is resistant to firmware-tampering attacks. Nevertheless, an attacker may be able to install a malicious impersonation—swapping out a smart meter, for example, for a rogue unit that looks like the legitimate meter but that covertly sends private energy-accounting information to a malicious Web site.

Therefore, even when secure boot is used, users might require additional assurance that a deployed product is actively running the known-good TCB. Toward that end, electronic systems that are connected to management networks can use remote attestation.

In one simple approach that any electronic system can implement, the system would use IKE/IPsec (Internet key exchange/Internet Protocol security) or SSL (Secure Sockets Layer) with public-key cryptography in mutual-authentication mode to establish a secure channel to a remote attestation server. Within those protocols, the private key representing the electronic system's identity "signs" the data that the attester subsequently authenticates. As long as the TCB validated during secure boot includes the private key and the client side of the secure-connection-protocol software, the attester

Figure 2 Secure boot requires verifying the chain of firmware components that execute up to and including the OS/hypervisor that ultimately controls the platform.

has assurance that the system is running the expected firmware.

Of course, a hardware root of trust is useless if it validates and launches buggy software (**Reference 2**). Electronics designers need a software root of trust—a trustworthy operating system—on which to develop and deploy trustworthy applications.

Monolithic operating systems contain system-software elements—such as networking stacks, file systems, and device drivers—that share a single memory space and execute in privileged mode. That structure results in a large TCB and thus provides a plethora of opportunities for hackers. Examples of monolithic operating systems include Windows, Linux, and VxWorks.

In contrast, a microkernel operating system runs only a minimal set of critical services—such as thread management, exception handling, and interprocess communications—in supervisor mode and hosts other system components in user mode. Such a setup permits access only to those resources the designer has deemed appropriate. A failure in one component does not harm other components managing independent resources. Because a microkernel is simpler than a monolithic operating system, the microkernel's security can be ensured and verified more easily (**Figure 3**).

Most general-purpose operating systems are monolithic because they were designed more than 20 years ago, when microprocessor performance limitations encouraged a monolithic approach. User applications can access most services with an efficient kernel system call.

In contrast, a microkernel implements services in communicating processes. For example, an application that seeks to access a file via NFS (the Network File System) may require communications among the application, a TCP/IP process, a device driver process, and an NFS process. This work is behind the scenes: An application calls the typical read() or write() interface, and the microkernel routes data among the system processes.

Microprocessor speed has rendered irrelevant the process/messaging overhead associated with microkernels. All major operating systems are now designed with microkernels and run practically every type of electronic product, including smartphones, avionics, networking gear, and process-control systems. Examples of microkernel-based operating systems include Integrity, LynxSecure, Neutrino, and PikeOS.

It is unlikely that Linux would be monolithic if created today. As Linux pioneer Linus Torvalds famously stated in a response to Andrew Tanenbaum's 1992 critique of the open-source operating system, "True, Linux is monolithic, and I agree that microkernels are nicer. From a theoretical (and aesthetical) standpoint, Linux loses" (**Reference 3**).

Nevertheless, Linux has an enormous and growing base of applications—and programmers—that designers want to leverage. The decision to use a microkernel or a monolithic operating system need not be mutually exclusive. Virtualization provides the ability to incorporate both.

SYSTEM VIRTUALIZATION

The motivations for system virtualization in the data center, such as resource optimization and improved service availability, are well-known. Virtualization technology has broader
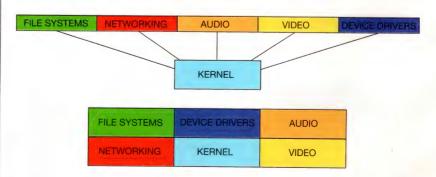

Figure 3 A microkernel OS (top) hosts system services in modular processes; a monolithic OS employs a single, large kernel.
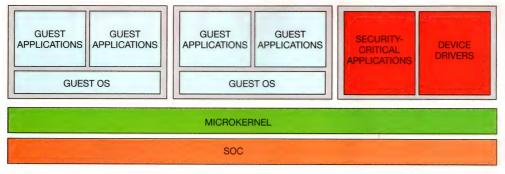
Figure 4 A microkernel-based hypervisor architecture enables both guest operating systems and lightweight applications that execute directly on the microkernel.

applications throughout the electronics world, however, and virtualization techniques play a critical role in security.
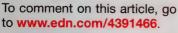
In 2005, Intel launched Intel VT (virtualization technology), which both simplified and accelerated virtualization. Recently, Intel VT has become available in Intel Atom embedded processors. Similar hardware assists for virtualization have emerged in other CPU architectures, including ARM (via ARM virtualization extensions), enabling efficient virtualization in all forms of electronics.

A microkernel-based hypervisor architecture includes computer virtualization as a service on top of a trustworthy microkernel (**Figure 4**). Examples of microkernel hypervisors include Green Hills Software's Integrity Multivisor and some variants of the L4 microkernel. Such products realize microkernel robustness while enabling reuse of general-purpose operating systems such as Linux. Electronics designers can make dramatic improvements in overall system security by applying this architecture intelligently, moving a small number of critical security functions outside of the legacy environment and hosting them directly on the microkernel TCB.

SECURE RETROFIT

In 2010, a widespread VxWorks remote-management vulnerability was discovered: An open system-diagnostics port

could enable hackers to install malware, even rooting or replacing the operating system itself. A basic defense is to guard the management port with mutual authentication using network-security protocols such as SSL.

An SSL connection, however, is the tip of the iceberg in ensuring secure connectivity. Without proper safeguards, including testing at the binary level and secure delivery, attackers can gain access using a range of proven techniques. Designers can employ a high-assurance development process for their own software, but how can they protect against vulnerabilities in third-party operating systems, many of which are shipped in binary form only and lack any security provenance, pedigree, or indemnification?

System virtualization can enable the retrofit of secure connectivity to electronic systems. The technique involves uplifting the legacy operating environment into a virtual machine, securely isolated from critical functions—such as connection authentication and configuration management—provided by a trustworthy hypervisor. The device management software can also be used to monitor, configure, and patch the legacy operating system itself (**Figure 5**).

DATA-STORAGE PROTECTION

In 2010, a CBS News segment demonstrated that office copiers were gold mines for private information that could easily be harvested from disk drives in the machines. The systems in question lacked data-storage protection. Many modern electronic systems, however, have encrypted storage requirements, driven by intellectual-property

protection, digital-rights management, and more.

With FDE (full disk encryption), the entire storage medium is encrypted, ensuring that hidden files, such as operating-system swap files, are not exposed. SEDs (self-encrypting drives) handle FDE within the storage peripheral itself, but because of form-factor limitations many electronic systems are unable to accommodate stand-alone SEDs. Encryption can be performed at the next level up: the device-management layer, typically a block-oriented driver. Protection at this level can still cover the entire managed device.

The storage encryption key must never be stored in the clear on the media, but it is often necessary to store an encrypted copy of the key. The key is unwrapped for active use while the system executes as authorized.

For interactive systems such as smartphones or point-of-sale devices accessed by a store clerk, successful user authentication is the trigger that unwraps the key. For unattended electronic systems, the symmetric key can be retrieved using a secure network connection to a remote server that holds a database of the provisioned data-encryption keys. The electronic system initiates key loading whenever a data-encryption key must be unlocked, such as at boot time.
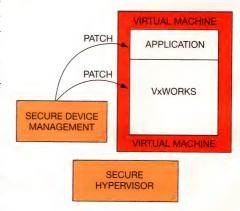

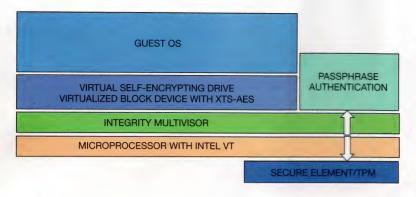Figure 5 A hypervisor can enable the retrofit of secure device management services to legacy systems.

Figure 6 A secure hypervisor enables a virtual self-encrypting drive.

VIRTUAL SEDs

Data protection is another emerging application for system virtualization in connected devices. Here, a hypervisor launches the main operating system in a virtual machine and executes the storage encryption components outside of the virtual machine, creating a vSED (virtual self-encrypting drive). All security-critical aspects of data-at-rest protection—authentication to unlock the SED, the encryption/decryption protocol, and key management—occur outside of the primary operating-system environment, protecting the stored data from external attack (**Figure 6**).

The hypervisor virtualizes the physical storage. When the operating system writes a block of data, a trusted application encrypts the block and stores it to an appropriate block of the physical device. Because encryption is executed in a securely partitioned application, no untrusted software on the platform can access the encryption key or perform cryptanalytic attacks.

On-chip cryptographic accelerators should be used for efficiency whenever possible. An alternative to a dedicated crypto coprocessor is to offload encryption on one core of a multicore processor, whose full complement of cores is rarely fully burdened.

When a passphrase is used to unlock storage, designers must be aware of key-logging and screen-scraping malware attacks in the primary operating-system environment. In the vSED environment, the passphrase input and associated graphical screen management (if applicable) are performed in native applications, protecting against such attacks.

Electronics developers who must design security into their products confront a plethora of design choices and constraints. Designers must take maximum advantage of silicon technologies that enable secure boot, efficient encryption, and secure isolation of data-protection subsystems using virtualization.**EDN**

REFERENCES

[1] Koscher, Karl, et al, "Experimental Security Analysis of a Modern Automobile," *Proceedings of IEEE Symposium on Security and Privacy,* IEEE Computer Society, May 2010.

[2] Kleidermacher, David, and Mike Kleidermacher, *Embedded Systems Security: Practical Methods for Safe and Secure Software Systems Development,* Elsevier, 2012.

[3] Kleidermacher, David, "Monolithic stupidity," *EETimes.com*, Nov 18, 2011, www.eetimes.com/4230750.

AUTHOR'S BIOGRAPHY
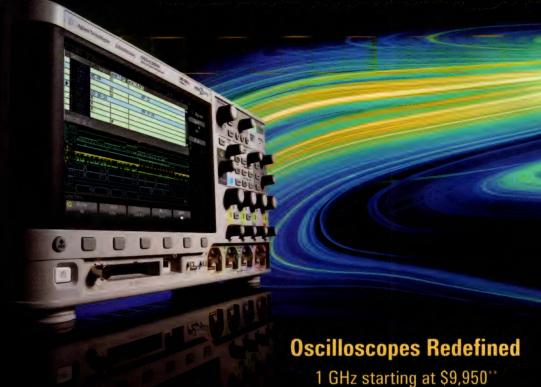
David Kleidermacher is chief technology officer at Green Hills Software, where he is responsible for technology strategy, platform planning, and solutions design. Kleidermacher is an authority in systems software and security, including secure operating systems, virtualization technology, and the application of high-robustness security-engineering principles to solve computing-infrastructure problems. He earned a bachelor of science degree in computer science from Cornell University (Ithaca, NY) and has been with Green Hills Software since 1991.

BY KEVIN KING · RENESAS ELECTRONICS AMERICA

Managing the 8- to 32-bit processor migration

TO PREVENT SEEMINGLY LILLIPUTIAN PROCESSOR DISPARITIES FROM HATCHING A SWARM OF BUGS, APPLICATION DEVELOPERS WORKING IN C MUST KEEP NATIVE PROCESSOR SIZE AND ENDIANNESS IN MIND AND WRITE THEIR CODE WITH AN EYE TOWARD PORTABILITY.

Back when I started in electronics, working on discrete, 4-bit processors, I couldn't have known I would one day have to worry about how big an integer was or discuss processors in a *Gulliver's Travels* context. As geometries shrank and prices dwindled, however, there was a great migration of applications from 8- to 16- and then to 32-bit processors. Along the way, tools evolved to bring code generation and application development to new levels of efficiency—generating more headaches in the process.

The problem had its genesis in the engineers working with the first microcontrollers who assumed 16 bits for an integer would be "good enough." Indeed, the early mainframe and minicomputer architectures differed in word length as well as in bit and byte ordering; the number of bits in an integer related to the architecture's word length and varied from machine to machine.

With apologies to Jonathan Swift (**Reference 1**), engineers have revised the Lilliputians' argument to debate which end of a number—the largest (big-endian) or smallest (little-endian)—should come first in memory. There are valid arguments on both sides of the "endianness," or byte-order, debate (**Reference 2**), but this article focuses on the ramifications for developing applications using C code.

An engineer embarking on a project using a 32-bit, little-endian processor, for example, might dismiss discussions of native processor size and endianness as merely academic for a new design. Then, once the project is under way, the engineer might discover that:

• The company has code intellectual property that is "tried and true" but has never been tested on a 32-bit, little-endian processor and must be recompiled from C source code;

• The code IP must be reusable by another internal group;

• The little-endian processor will pass data structures over the communications link to a big-endian processor; and

• The project must also use a 16-bit, big-endian memory-mapped device that the company's ASIC group has provided.

Any of those scenarios can cause problems if the engineer does not keep in mind the native size and endianness of the processor.

VARIABLE SIZE

Consider what happens when our hypothetical engineer is given code that appears to be a very simple structure for the calendar portion of a real-time clock:

```
union date_union{
  unsigned long year_month_day;
  struct{
    unsigned int  year;
    unsigned char month;
    unsigned char day;
  }word_and_byte;
} date = { 0x20121022};

static unsigned char read_day;
static unsigned int read_year;

  read_month = date.word_and_byte.month;
  read_day = date.word_and_byte.day;
  read_year = date.word_and_byte.year;
```

The code was built without yielding errors, but neither does it yield the expected results. The structure is so simple; what could have gone wrong? The engineer looks at the variables in the watch window, which shows the following:

```
read_month    H'00 '.'
read_day      H'00 '.'
read_year     H'20121022
```

It turns out the code had only been used on 16-bit processors up to this point. The first problem is that ANSI C does not define the number of bits in an integer; that number is typically related to the native size of the accumulator in the processor.

Thus, the first point is to write your code so that the size of the variables is known. The C99 standard addresses this issue through fixed-width integer support (**Reference 5**). So, our hypothetical engineer plugs in his solution to correct the variable size (he happens to use the C99 conventions, so if his compiler were C99-compliant he could have just included <stdint.h> in his file):

```
#define uint32_t unsigned long
#define uint16_t unsigned short
#define uint8_t  unsigned char
```

ENDIANNESS AND PORTABILITY

The engineer then uses those definitions throughout the "tested" IP and runs the code. Now the watch window, which should show Oct 22, 2012, instead shows the following:

```
read_month  H'12 '.' { 00001401 }
   (unsigned char) [Current Scope]
read_day    H'20 ' ' { 00001400 }
   (unsigned char) [Current Scope]
read_year   H'1022 { 0000147E }
   (unsigned short) [Current Scope]
```

The 16-bit processors on which the code had been used earlier were big-endian; thus, the initialization string is being put into memory in the wrong order. Our engineer, bitten once, revisits the four bullet points and vows to solve all his problems by supporting both big- and little-endian structures and making the code portable.

That decision, however, has its own implications: Efforts to make code portable between machines of different endianness can uncover further, more obscure issues.

A common practice when coding is dereferencing with a pointer and casting at the same time. Our engineer writes some new code and passes it to a colleague on another project (recall the second bullet point on pg 49); he remembers to use his new definitions, in case the other project is using a 16-bit machine:

```
uint16_t short_data = {0x1234};
uint8_t  *p_char;
uint16_t s_data;
uint8_t  c_data;

   s_data = short_data;
   p_char = (unsigned char *)&short_data;
   c_data = *p_char;
```

The colleague from the other group soon comes back in a huff and tells the engineer the code doesn't work. The watch window shows a value of 0×12 for c_data, when clearly the lower byte of the data is 0×34.

The problem now is that the combination of casting and dereferencing causes the pointer to look at the wrong byte when used on a big-endian processor. The take-away is to use casting with caution, especially when using pointers to reference the data. There are a number of ways to skin this cat; **Listing 1** shows our engineer's solution.

EXTERNAL BUS ACCESS

As for the fourth bullet point on pg 49—accommodating a big-endian ASIC—the easy solution is to use an MCU whose external controller offers big-endian support. The MCU will then use big-endian byte ordering to the ASIC during a chip select, read, or write and will look for the data on the correct byte lane.

If the bus does not support big-endian structures, a direct access would put the data in the wrong register in the ASIC (**Figure 1**). One solution would be simply to reorder the data before writing it out, much as bytes might be swapped to satisfy a big-endian protocol in communications. That approach, however, risks the mistake of putting the data on the incorrect byte lane for a byte write. By contrast, if the selected processor has big-endian programmability on the external bus, the hardware handles all byte reordering, including the proper byte decode for byte accesses (**Figure 2**).
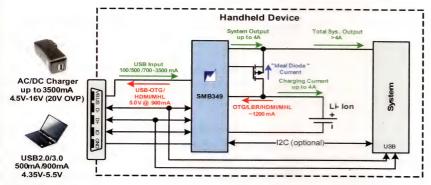
Bit-field operations are a powerful tool. Big/little-endian ordering typically extends down to the bit level; that is, a bit

Finally! Charge Big Batteries Fast With 4A Chip-Scale Li-Ion Charger ICs

Flexible USB/AC Input from +3.6V to +16V, Tiny Solution Size, I²C Programmable Parameters and Built-In Safety – Ideal for Tablets and Smartphones
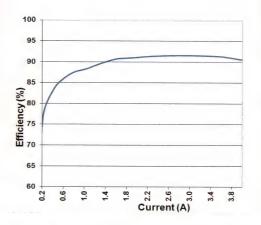

Applications
- Tablets
- Smartphones
- E-Readers
- UltraBooks
- Battery "JuicePacks"
- Portable Gaming
- Portable Digital Video

Features
- +3.6V to +16V Operating Input (+20V OV Protection)
- Fast-Charging, Flexible 4A Switch-Mode Architecture
  - *TurboCharge™** current-multiplier cuts charge time by 30%-60%
  - *TurboCharge+™** auto-float voltage control (AFVC) further reduces charge time by compensating for internal battery impedances
  - *CurrentPath™* with dual outputs for system/battery (SMB349) supports instant-on with dead/missing battery
  - *FlexCharge/FlexCharge+™** auto power source detection (APSD/AIVD) per USB2.0/3.0/BC1.2 to detect USB or AC/DC source +5V to +16V
  - *OptiCharge™** auto input current limit (AICL) detects and adapts to source current limit to maximize available power
- I²C Programmable Parameters and Functions with NV Configuration
- *SafeCharge™* safety features support JEITA/IEEE1725
  - Battery and IC over-voltage/current/temperature protection
  - Trickle charge for deeply-discharged cells
  - Safety timers and fault monitors/reporting
- Tiny 3.2 x 3.0 CSP and 52mm²/1.2mm z-height solution size

Industry's Highest System Efficiency

	SMB349	SMB359	SMB347	SMB137C	SMB346	SMB136C
Input Voltage Range (V)*	4.35 to 16 (20)	4.35 to 16 (20)	4.35 to 6.2 (20)	4.35 to 6.0 (18)	4.35 to 6.2 (20)	4.35 to 6.0 (18)
# of Inputs/Outputs	1/2	1/1	2/2	2/2	2/2	1/2
Maximum Charge Current (mA)	4000	4000	2500	1500	1250	1500
Maximum Input Current (mA)	3500	3500	2500	1500	2500	1500
CurrentPath™ Control	√		√	√	√	√
Charge Current Voltage Output	√	√	√		√	
Low-Battery Recovery Mode				√		√
Automatic Power Source Detection **	rev 1.2	rev 1.2	rev 1.1/1.2	rev 1.2	rev 1.1/1.2	rev 1.2
Package	3.2x3.0 CSP-49 5x5 QFN-40	3.2x3.0 CSP-49 5x5 QFN-40	3.0x2.5 CSP-30	3.0x2.5 CSP-30	3.0x2.5 CSP-30	3.0x2.5 CSP-30
Solution Size (mm2)	52	52	32	38	32	35

All chargers have Battery Thermal Protection & JEITA Support, IC Thermal Protection, Auto Input Current Limit, Safety or Watchdog Timers, Programmable Charging Parameters, I2C Interface, USB On-The-Go, TurboCharge™ Mode*

* Patent granted or pending

For more information see:
www.summitmicro.com/SMB349

SUMMIT
MICROELECTRONICS
Programmable Power for a Green Planet™

field can be left- or right-justified within a word. Little-endian processors typically default to right-justified bit fields, allocating bits from the lower (little) end first; big-endian processors are left-justified, allocating bits from the upper (big) end first.

Figure 1 If the bus does not support big-endian access, a direct access would put the data in the wrong register in the ASIC.
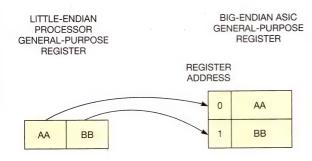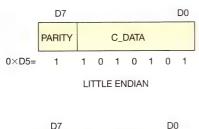

Figure 2 If the chosen processor has big-endian programmability on the external bus, the hardware handles all "byte reordering," including the proper byte decode for byte accesses.
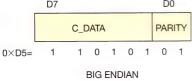

Figure 3 As soon as code that is not portable is moved to a big-endian processor, it fails to operate as expected; this "incompatibility" is easy to see in the figure, which show the bits in memory with a value.

Consider the following code, which is not portable:

```
typedef union char_union{
  uint8_t char_data;
  struct{
    uint8_t c_data:7;
    uint8_t parity:1;
  }BIT;
} comm_char;

uint8_t proc_comm_char(comm_char * my_
  com_char)
{
  return(my_com_char->BIT.c_data);
  // strip off the parity bit
}
```

As soon as the code is moved to a big-endian processor, it fails to operate in the expected manner. The "incompatibility" is a little easier to see in **Figure 3**, which shows the bits in memory with a value. The parity bit is at 0 (MSBit first in memory), and the char is not aligned to the byte; it was on bit 7 in the little-endian processor. Thus, for any declared structure containing bit fields, be wary of any bit manipulations done on a whole word (& or |) unless such manipulation is required, as in the case of some SFRs (special-function registers) in the MCU.

Again, once you are cognizant of the system, the fix is trivial, though avoiding the issue in the first place is even better. Our engineer's fix looks like this:

```
typedef union char_union{
  uint8_t char_data;
  struct{
#ifdef LITTLE_ENDIAN
    uint8_t c_data:7;
    uint8_t parity:1;
#else
    uint8_t parity:1;
    uint8_t c_data:7;
#endif
  }BIT;
} comm_char;
```

Note that some compilers and processors let you set the bit order, left or right, when defining bit structures.

NATIVE CORE OPERATIONS
The native size of the processor is another consideration. When building embedded devices (the embedded Linux crowd excepted), engineers are concerned about "optimization" of ROM size, RAM size, and sometimes even logic size. It is common practice among MCU designers to align SFRs or peripheral registers with specific buses to reduce the metal and the bus switches required to get the registers into the correct part of the general register. For example, the upper byte of a 32-bit register as a character requires the device to steer the byte from the upper bus to the lower bus while presenting a 32-bit aligned address; otherwise, it might read the byte presenting an "unaligned" address to a chip that does not support reading the peripheral in this way.

Be cognizant of your SFR/peripheral register alignment and access rules. If the hardware manual for the device tells you to read the SFR a certain way, be sure to code that way. Some compilers take care of this with #pragmas; others create specific iodefine files for the MCU that will access the peripheral a certain way if you use the structures or a combination of the two methods.

MEMORY ACCESS

We like to think of RAM as being universally addressable no matter how we write our code. When using an 8-bit processor, regardless of how we wrote our code, the fundamental access to the memory was 8 bits at a time. In the era of 16- and 32-bit units, however, the generalization no longer applies. Revisiting the calendar-structure code sample used previously, but this time with a slightly different syntax, illustrates the point:

```
#pragma pack
union date_union{
  uint32_t year_month_day;
  struct{
    uint8_t  day;
    uint16_t year;
    uint8_t  month;
  }word_and_byte;
} date;
#pragma unpack
```

I would describe the code sample as less than portable. The engineer has packed the structure—that is, there are no padding bytes to align word and long variables—and since a char variable comes first in the structure, it causes the 16-bit year variable to be "unaligned" (16 bits on a byte address). That's typically fine if your accesses are made via a direct reference to the structure; the compiler can "fix" things. Consider what happens, however, if the reference is made by a pointer:

```
p_my_year = &date.word_and_byte.year;
current_year = *p_my_year;
```

The result will be a 16-bit access to the unaligned location of date.word_and_byte.year. That's fine for processors that have hardware in their bus state controllers to split the data into multiple byte reads and reconstruct the 16-bit word. But for many processors—particularly DSP and RISC processors for which high speed is key—unaligned access will cause an illegal address exception. Even processors that allow unaligned access exact a small performance penalty for the split transaction. Armed with that understanding, you can make an informed decision between a slight performance hit for multiple accesses or a slight RAM hit for not packing the structure, and you can avoid this kind of situation when building code for processors that don't support "unaligned" accesses.

Remember this old trick when declaring structures: If you declare from the largest element first down to the smallest

EVEN PROCESSORS THAT ALLOW UNALIGNED ACCESS WILL EXACT A SMALL PERFORMANCE PENALTY FOR THE SPLIT TRANSACTION.

element, with most compilers, unaligned accesses typically will not occur even when using pointers. That's because the compiler/linker will typically align on the boundary of the first element, and the remaining elements will end up on the boundary that is correct for their size.

MATH OPERATIONS

When migrating from smaller to larger processors, math operations are typically not an issue. The 32-bit processors have barrel shifters and MAC units, and some even have floating-point units, so capability is almost a given. But remember our engineer's first bullet point: He would be using tried-and-true code. What might he encounter, and what would be the resultant considerations?

One common problem is looking for overflow on integers. Using fixed-width integers as defined in C99 can go a long way toward solving the problem; as already noted, however, given the decreasing price of 32-bit MCUs and the demand for high functionality in embedded devices, engineers are migrating to them in ever-increasing numbers. As a result, they are seeing code such as the following:

```
int my_multiplier;
int my_int;

  my_int *= my_multiplier;
    /* calculate the new result */
  if (my_int & 0x8000)
    /* check for overflow */
  {
    overflow_error();
      /* flag the error */
  }
```

Many 8-bit MCUs do not do so well when accessing larger variables, so you might even see this written as an access to the upper byte only to save cycles. The bottom line is to be wary of any "hard coded" values that are used to "manipulate" numerical variables as they relate to sign or overflow operations.

When developing code for the first time or porting code from an 8- to 32-bit processor, issues believed to be trivial can prove otherwise. Keeping the size of the processor in mind, planning ahead, and writing your code with an eye toward portability should help you avoid such situations.

By the way, which end of the egg do you break?EDN

REFERENCES

[1] Wikipedia contributors, "Lilliput and Blefuscu," *Wikipedia, The Free Encyclopedia*, July 18, 2012, http://en.wikipedia.org/wiki/Lilliput_and_Blefuscu.
[2] Wikipedia contributors, "Endianness," *Wikipedia, The Free Encyclopedia*, July 18, 2012, http://en.wikipedia.org/wiki/Endianness.
[3] Netrino, Embedded C Coding Standard, 2009, www.netrino.com.
[4] Kernighan, Brian, and Dennis Ritchie, *The C Programming Language*, 2nd Ed, Prentice Hall, 1988.
[5] International Organization for Standardization and International Electrotechnical Commission, ISO/IEC 9899:1999(E), Second Edition, 1999-12-01, ISO/IEC 1999, ANSI Approval date 5/22/2000, www.iso.org.
[6] PC-lint for C/C++, www.gimpel.com/html/index.htm.
[7] Misra C:2004, "Guidelines for the use of C language in Critical Systems," www.misra.org.uk.

AUTHOR'S BIOGRAPHY

Kevin King is a senior staff applications engineer with Renesas Electronics America Inc. He is currently working with a global development team on an optimized DSP library for the RX600 family of microcontrollers. King has published application notes in the fields of medical devices, motor control, and digital filtering.

RECOMMENDATIONS

Approaches for the 32-bit migration include, but are not limited to, these recommended actions:

• Define and use a coding standard within your company;

• Use static analysis, such as PC-lint (Reference 6) and Misra C (Reference 7) checkers when possible and appropriate;

• Understand variable size and reduce its effect on portability when writing your code by forcing to a known size (C99 <stdint.h> helps in this regard);

• As a minimum, when coding style alone does not yield portability, use defines such as BIG_ENDIAN and LITTLE_ENDIAN to make your code build correctly for any MCU type your company might choose (your compiler may have a predefine);

• If possible, code your functions for portability (doing so may require a slight increase in code size) by using initialization routines not tied to the endianness of the MCU;

• Understand the limitations of your processor when it comes to accessing different-sized units within its special-function register and memory spaces (bit, unaligned, and packed structures);

• Understand your application from a macro rather than a micro perspective (to whom or what will your application "talk"?); and

• Understand the portability requirements of the code you are using or writing, and be aware of any casting and pointer dereferencing that might get you into trouble.

Note: The use of #ifdef LITTLE_ENDIAN here is for illustration only; you may choose some other designation or identifier. Many processors that support big-endian data will have a predefined preprocessor identifier.

designideas

Quasistatic Spice model targets ceramic capacitors with Y5V dielectric

Hugo Coolens, KaHo Sint-Lieven, Ghent, Belgium

Ceramic Y5V SMT capacitors have recently become available in values and sizes that were previously available only with electrolytic capacitors. At first glance, they may seem a worthwhile unpolarized alternative to electrolytic capacitors, and they sometimes are. These capacitors, however, have a capacitance that is a function of the applied voltage. Modeling them as ordinary linear capacitors can lead to great discrepancies between simulated and measured—let alone expected—results.

Figure 1 shows some measurements of capacitance as a function of dc voltage on a 10-µF, 25V Y5V 1206 capacitor. **Reference 1** describes a method of modeling nonlinear capacitors using a look-up table. The use of a look-up table, however, adds complexity to the simulation and can lead to conver-

gence problems. If you limit the operating voltage to approximately 80% of the capacitor's voltage rating, you may find that a simpler quadratic quasistatic model can be sufficient for quickly arriving at a first approximation.

Fitting the measured data of C to $a_0 + (a_1 V_{DC}) + (a_2 V^2_{DC})$ using quadratic regression yields the coefficients of 8.500065×10^{-6}, -7.445791×10^{-7}, and 1.922001×10^{-8} for a_0, a_1, and a_2, respectively. **Listing 1** shows the conversion from the parameters to the nonlinear-capacitor model in PSpice. You can see the equivalent capacity in PSpice's graphical postprocessor Probe as i(c)/(2*3.14159*frequency).

The differences between the measured values and the quadratic approximation are at worst approximately 20% (**Figure 2**). The relative residues are normally distributed with a mean of

2.5% and a standard deviation of 10%. Those deviations might seem to be a rough approximation for a model, but you must compare them with the normal linear C model, for which the relative
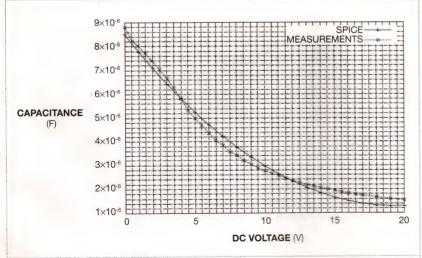
DIs Inside

▶To see and comment on all of *EDN*'s Design Ideas, visit www.edn.com/designideas.

▶Submit your own Design Idea to edndesignideas@ubm.com.

Figure 1 This graph measures capacitance as a function of dc voltage on a 10-µF, 25V Y5V 1206 capacitor.

LISTING 1 CONVERSION PARAMETERS

```
quadratic y5v–cap model
. param a0=8.500065u
. param a1=−744.5791n
. param a2=19.22001n
. param valdc=1
c 1 0 cmod {a0}
vin 1 0 dc { valdc } ac 1
. model cmod cap
   ( c=1 vc1={a1/a0 }
     vc2={a2 /a0 })
. ac dec 20 100 100k
. step lin param valdc 0 25 1
*display equivalent c as
   i ( c )/(2*3.141 5 9 *frequency)
. probe
. end
```
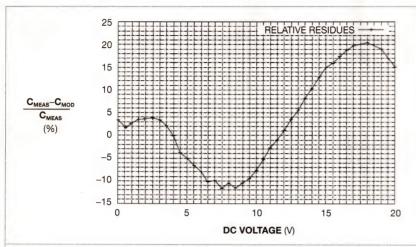
designideas

Figure 2 The differences between the measured values and the quadratic approximation are at worst approximately 20%.

The y-axis is labeled $\dfrac{C_{MEAS}-C_{MOD}}{C_{MEAS}}$ (%)

approximately 5 or 6V. That result is not surprising, but it again shows that you can't just replace an electrolytic capacitor with a ceramic one of the same capacitance and voltage rating.EDN

REFERENCES

[1] "A Nonlinear Capacitor Model for Use in the PSpice Environment," Cadence Design Systems, Dec 2009, http://bit.ly/J5RGc6.

[2] Cain, Jeffrey, PhD, "Comparison of multilayer ceramic and tantalum capacitors," AVX Corp, http://bit.ly/RuwKP5.

[3] "LT3083 Adjustable 3A Single Resistor Low Dropout Regulator," Linear Technology Corp, pg 12, http://bit.ly/JCnT88.
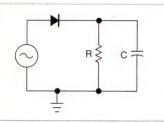
Figure 3 A simple rectifier setup checks the validity of the model.

error at 50% of the dc-voltage rating is already approximately 300%. Note that in practical applications you must limit the capacitor terminal voltage to no more than 80% of the capacitor rating so that the model almost always gives a better approximation than the worst-case value might suggest.

A simple rectifier setup checks the validity of the model (**Figure 3**). It measures a peak-to-peak ripple voltage of 10.2V (**Figure 4**). Simulating the circuit with the quadratic model in PSpice shows a ripple voltage of 10.4V (**Figure 5**). You don't need PSpice for this model; Spice 2G6 also has a built-in feature for modeling this type of capacitor. You enter it in this format:

```
c1 2 0 poly 8.500065u
 -744.5791n 19.22001n
```

PSpice no longer has this feature, even though it is said to be Spice 2G6-based.

Note that replacing the capacitor with an ordinary electrolytic that has the same capacitance and voltage rating would reduce the ripple voltage to
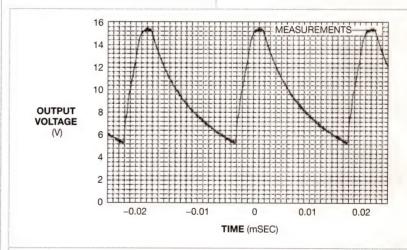

Figure 4 The setup measures a peak-to-peak ripple voltage of 10.2V.
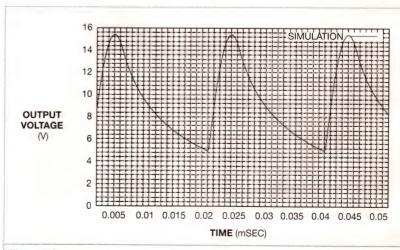

Figure 5 Simulating the circuit with the quadratic model in PSpice shows a ripple voltage of 10.4V.

Securing Energy

Proven security solutions protect the grid's endpoints

Attacks to the smart grid can occur from any point on the network. From stealing electricity and customer data to threatening critical assets, cyber criminals can deploy threats to the grid miles from the control room.

Maxim's Newport smart meter development platform protects the endpoints of the grid with advanced cryptography, physical-attack-detection technology, and sensitive data protection so the action doesn't stop.

Proven Security

- Accelerators support asymmetric schemes including RSA, DSA, and elliptic curve
- True hardware random number generation generates stronger challenges and keys
- Tamper inputs detect physical attacks
- Built-in IP protection for supply chain and manufacturing process

Get Started with Built-In Security

Newport Development Platform

Free whitepaper: "Securing Smart Meters."
www.maxim-ic.com/newport

MAXIM
INNOVATION DELIVERED®

designideas

Microcontroller drives piezoelectric buzzer at high voltage through one pin

Mehmet Efe Ozbek, PhD, Atilim University, Ankara, Turkey

↘ A previous Design Idea demonstrates how you can use a microcontroller to drive a piezoelectric buzzer at a high alternating voltage through a four-MOSFET circuit that interfaces to two of its I/O pins (**Reference 1**). This expanded Design Idea provides a modification of the previous circuit to save one of the I/O pins of the microcontroller. Q_4's gate connects to Q_2's drain rather than a second I/O pin (**Figure 1**). The microcontroller turns on Q_2 by applying a high logic level to the I/O pin, pulling Node A down to a low logic level. This action turns on Q_3 and turns off Q_4. The voltage on Node B becomes 15V, and Q_1 turns off. The voltage across the piezoelectric element is now 15V.

The microcontroller then toggles the I/O pin low, turning off Q_2. Q_1 is also off, so Node A slowly rises to a high logic level through pullup resistor R_1. When the voltage on Node A reaches the switching threshold of the inverter comprising the Q_3 and Q_4 pair, Q_3 quickly turns off and Q_4 quickly turns on. The consequently low logic level on Node B turns on Q_1 and speeds the increase of Node A's voltage. The 15V across the piezoelectric buzzer is now of the opposite polarity.

R_2 weakens the coupling between the output and the input of Q_4 due to the presence of the piezoelectric element. A value of 330Ω for R_2 is usually sufficient to suppress high-frequency oscillations that the feedback causes. The drained power from the supply increases if you use low values for R_1. Using excessively large values for R_1 also increases power dissipation by prolonging the switching of the transistors and associated shoot-through currents. The optimum value for R_1 is approximately 1 $k\Omega$.

Saving an I/O pin with this design involves the trade-off of increased power consumption. The circuit's power consumption is thus one order of magnitude greater than the circuit described in the previous Design Idea.**EDN**

Figure 1 One microprocessor I/O pin drives this circuit to generate an alternating voltage across the piezoelectric buzzer.

REFERENCE

1 Ozbek, Mehmet Efe, "Microcontroller drives piezoelectric buzzer at high voltage," *EDN*, March 1, 2012, pg 44, http://bit.ly/JyzLpz.
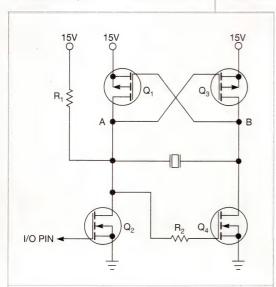

Circuit automatically switches off DMM

Vladimir Oleynik, Moscow, Russia

↘ The circuit in **Figure 1** lets you install this automatic-shutoff function in a DMM (digital multimeter). When the momentary pushbutton switch is open, the capacitor discharges through the resistor. The transistor and DMM are off because the gate-to-source voltage is 0V. When you momentarily press the button, the capacitor immediately charges up to the battery voltage. The transistor's gate voltage is higher than its source, turning on the DMM. When you release the button, the capacitor begins to discharge slowly through the resistor. When the gate voltage reaches the threshold level, the transistor turns off, thus switching off the DMM.

With a fresh battery and the values in **Figure 1**, you have about 50 seconds to test the circuit. Of course, you can set this time according to your requirements by changing the RC-network values. If your DMM starts switching on for a shorter time after you press the button, the battery is reach-
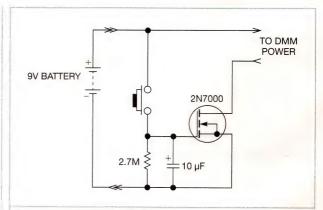

Figure 1 Because of the small number of components required, this auto-shutoff circuit is easy to build and install inside the limited space of a DMM.

designideas

ing the end of its capacity and needs replacement.

Select a miniature, normally open pushbutton switch that will fit into a hole drilled into the circuit's DMM front panel. Because the DMM typi-cally does not have much free space inside, you may need to use minia-ture surface-mounted components on a small prototype board to minimize the circuit's size. Attach the board to the DMM's back cover using glue or a piece of double-sided adhesive tape. In some cases, through-hole components may do the job as well. To avoid acci-dentally turning on your DMM during transit, just turn its range switch to the off position.**EDN**

Reconstruct the input current in a grounded-impedance current sensor

Abhirup Lahiri, New Delhi, India

You can reconstruct the small-signal current flowing into a grounded impedance, which is part of a larger circuit. This current, such as the highpass-filter output for several current-mode biquadratic filters flowing into a capacitor—and which is a work-ing impedance of the filter—can be the summation of multiple currents.

Reconstruction of the output current by means of individually duplicating all constituent input currents through cur-rent mirrors can be error-prone, as the following **equations** show: $I_{IN}=I_1+I_2+I_3 ... +I_N$, and $I_O=\alpha_1 I_1+\alpha_2 I_2+\alpha_3 I_3 ... +\alpha_N I_N$. The current-mirroring coefficients, α_i, should ideally achieve unity to recon-struct I_{IN} accurately. All of these coef-ficients, however, not only deviate from their ideal values of unity but also can differ from each other because of current-mirror mismatches, including systematic mismatch due to the mir-ror's finite output resistance and random threshold mismatch for CMOS current mirrors. Those mismatches can cause the frequency character of the output's small-signal current—for example, a highpass response—to change. A useful technique is to sense the voltage across the impedance and then do a single-pre-cision voltage-to-current conversion, thereby getting rid of the errors arising from multiple mirroring operations.

Figure 1 shows a small-signal cur-rent-sensing circuit that must sense the total current flowing into an impedance, Z_1. The voltage-to-current conversion requires a matched load impedance, Z_2, and is built around an operational cur-rent conveyor (**Reference 1**). Any mis-match between the loads will cause the output current to be a scaled value of the input current, where the scaling factor is the ratio of impedances; hence, the frequency character of the small-signal output current will remain unchanged.

CURRENT-MIRROR MISMATCHES CAN CHANGE THE FREQUENCY CHARACTER OF THE OUTPUT'S SMALL-SIGNAL CURRENT.

The mismatch in either of the cur-rent mirrors is also of interest: Assuming that the input current of Current Mirror 1 is I_B+I_{IN} and the input current of Current Mirror 2 is I_B, then the output current is $I_O=\beta_1 I_B+\beta_1 I_{IN}-\beta_2 I_B=(\beta_1-\beta_2) I_B+\beta_1 I_{IN}$, where I_B is the quiescent cur-rent of the last stage of the amplifier and β_1 and β_2 are current-mirror mismatch-es. The output current has a dc-offset current, which you can easily cancel out, and a scaled value of the input's small-signal current, I_{IN}.**EDN**

REFERENCES
[1] Gift, Stephan JG, "Hybrid current conveyor-operational amplifier circuit," *International Journal of Electronics*, Volume 88, No. 12, 2001, pg 1225, http://bit.ly/L2CAle.
[2] Robinson, John, "New CCII Current Conveyor," Application Note 4198, Maxim Integrated Products, March 27, 2008, http://bit.ly/JVQ6rJ.
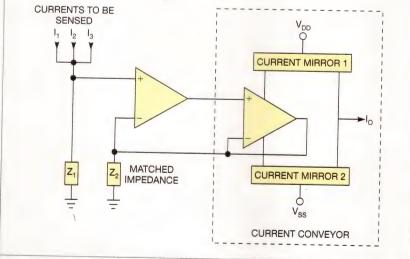
Figure 1 The current mirrors reflect the current draw of the load, Z_2, whose impedance—and, thus, its current—matches that of the impedance under test, Z_1.

designideas
CLASSICS

Originally published in the August 22, 1985, issue of *EDN*

Simple circuit suits quadrature detection

SL Black and HL Maddox, AT&T Technologies, Columbus, OH

The circuit in **Figure 1** generates an output voltage that you can measure to determine whether two sine waves have a quadrature relationship. If the output voltage is 0V, the inputs (ϕ_1 and ϕ_2) are exactly in quadrature. If the inputs are other than 90° out of phase, a dc voltage appears at the circuit's output. The voltage is proportional to the number of degrees that the input signals are out of quadrature. The polarity of the voltage is positive for phase angles of less than 90° and negative for angles of greater than 90°.

The signals **A** and **B** in **Figure 2** are in quadrature. When **A**'s signal is applied to the ϕ_1 input, a bilateral CMOS switch turns on during the positive half cycle and turns off during the negative half cycle. If **B**'s signal is applied to ϕ_2 simultaneously, an output similar to that of **C** appears at pin 2. Note that the areas above and below ground are equal. The integrating network, R_5C_1 in **Figure 1**, produces a net voltage of 0V.

If the phase angle is >90°, the area above ground is larger than the area below ground, and the output voltage is positive (**D**). If the phase angle is <90°, a negative output voltage results (**E**). If the 4016 triggers at a value other than 0V, the detector's accuracy will not change.

R_3, D_1, and D_2 provide input protection for the IC. The performance of the $R_4/R_5/C_1$ integrator depends on the frequency of the input signals and the impedance of the network at pin 1. If you choose 8.2 kΩ for R_1 and 2.2 kΩ for R_2, the values 8.2Ω, 4.7 kΩ, and 3.2 µF for R_4, R_5, and C_1, respectively, yield good performance at 25 kHz. These values will accommodate a 24V p-p swing at the ϕ_2 input. The values of V_{DD} and V_{SS} must be large enough to accommodate the input swings at the 4016. For example, an input swing of ±3V would call for 5V for V_{DD} and −5V for V_{SS}. **EDN**
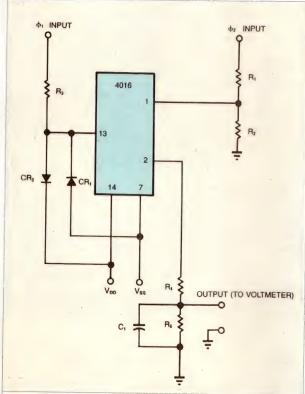

Figure 1 The bilateral switch in this circuit allows you to determine whether two sine waves are in quadrature. If the output voltage is 0V, the inputs (ϕ_1 and ϕ_2) are exactly in quadrature. If the output voltage is positive or negative, the waves are out of quadrature.
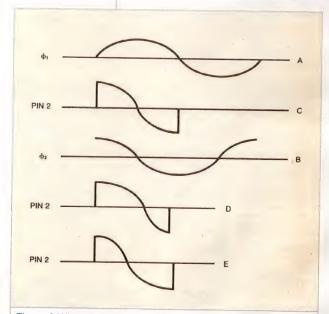

Figure 2 When ϕ_1 and ϕ_2 are in quadrature, the output of pin 2 (**C**) manifests equal areas above and below ground, resulting in a 0V integrated output from pin 2. If the waves are out of quadrature, a positive (**D**) or negative (**E**) voltage appears.

supplychain

LINKING DESIGN AND RESOURCES

The limits of lean

Ten years ago, "lean"—the practice of eliminating waste and non-value-added activities in the business environment—was the darling of the electronics industry. Like the "total quality management" revolution of the 1970s and 1980s, lean spurred an ecosystem of research, consultants, practitioners, awards, and benchmarks. And, like the practices of JIT (just in time) and build to order, lean reduced the levels of physical inventory in the electronics supply chain.

There's no question that lean achieves results. Wall Street analysts who track the electronics industry use low levels of inventory as a measure of sound financial management, hammering companies that hold too much on the shelf. In fact, until two natural disasters rocked the industry last year, nobody questioned whether lean was the right practice for the high-tech supply chain.

In March 2011, an earthquake and tsunami devastated parts of Japan and shuttered wafer-manufacturing facilities. In October 2011, flooding paralyzed Thailand, the global center of HDD (hard-disk drive) supply. In both cases, supplies of key electronics products were put in jeopardy because years of lean practices had eliminated inventory redundancies in the supply chain.

Right after the Japanese disaster, Malcolm Penn, founder and CEO of research firm and consultancy Future Horizons,

wrote on *EBN* that the JIT, on-demand, lean, batch, and outsourced-manufacturing models had taken over from inventory, work-in-progress, production lines, and multiple sourcing—despite the fact that the entire manufacturing process, from wafer-build to end-product delivery, takes six months ("Tech needs a healthier supply chain, part 1," http://bit.ly/N8Z8Q8).

In an early indication that the supply chain had been squeezed too far, "Nissan was forced to shut down car production due to a lack of engine management modules, itself the result of a shortage of ICs, despite the chip supplier's claim that it had made all the ICs originally asked for," Penn wrote. Future Horizons forecast an increase in such incidents.

Ultimately, there was no widespread post-tsunami shortage, but that was more a coincidence than a planned strategy. Toward the end of 2010, semiconductor inventories had begun building to what one researcher called "alarming" levels, and the overstock cushioned the supply-chain impact of Japan's natural disasters.

The industry wasn't so lucky in Thailand, where the clustering of HDD plants in a centralized location had eliminated redundancies to such an extent that roughly 70% of the world's HDD production was affected when the floods hit. HDD production is just now recovering.

The disk-drive industry is clearly reassessing its clustering strategy. Seagate Tech-

nology, for one, is reducing its number of inventory-holding JIT hubs in favor of value-added fulfillment centers closer to where end customers consume products. Dennis Omanoff, Seagate senior vice president for supply chain and procurement, noted in comments to his July 17 *EBN* blog entry ("Opportunities beckon as risks rise," http://bit.ly/QoJ0fi) that it's important to have strategic partners with global capabilities and regional locations to service the requirements of a particular geography. "This provides a more flexible network that can adapt rapidly to change while improving agility,

> ❝ Conversation around the supply chain is shifting to such adjectives as 'resilient' and 'agile.'

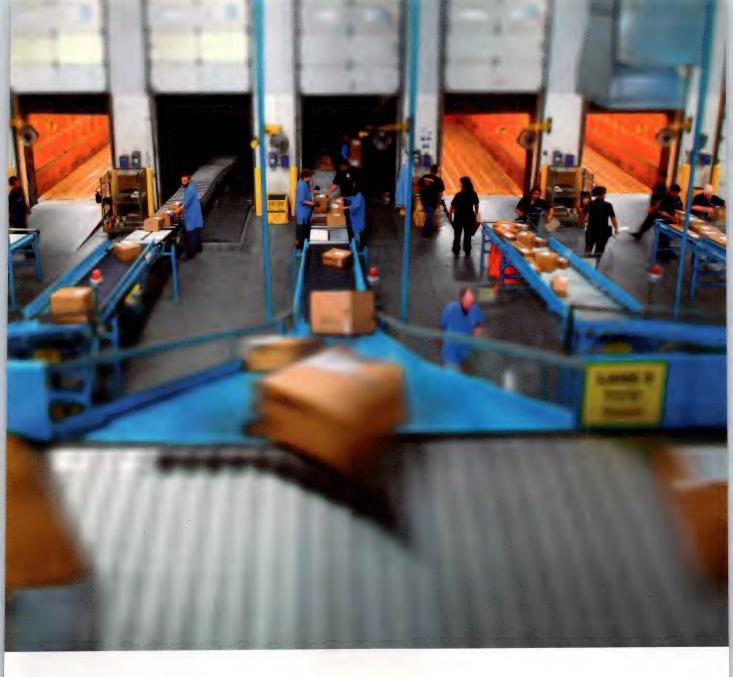responsiveness, velocity, and customer service," he wrote.

Since the 2011 disasters, conversation around the supply chain has been shifting from lean toward such adjectives as "resilient" and "agile." Gartner, in its annual analysis of leading supply-chain companies, noted that global companies are at an inflection point: "The past year brought global-scale supply-chain disruptions that [affected] multiple industries … These disruptions have even called into question whether supply chains have become too lean, requiring a fundamental change in approach" (http://bit.ly/MpVrH4).

It's unclear from Gartner and many other sources, however, exactly how to achieve

resiliency and agility. End-to-end visibility among partners is one of the components of a resilient supply chain, and an interesting thing in the electronics industry occurred after the Japan quake and tsunami: Rather than panic buy, customers called distributors to ensure that the components they'd ordered were actually on the shelf. They weren't interested in forecasts or whether orders were in process; they wanted to know where their physical inventory was. That approach runs counter to some of the principles of lean.

Several other trends indicate the supply chain is moving toward a middle ground between lean and gluttonous. In distribution, local sales and support offices are supplementing centralized hubs. Proximity warehouses are springing up closer to customers. Distributors at times even take advantage of opportunistic purchases to pad their inventory.

Lean has definitely increased the efficiency of the supply chain and has rendered many companies financially strong. But experts continually call for a reassessment of supply-chain strategies, and lean doesn't come up in those conversations as much as it used to.

—by Barbara Jorgensen,
EBN community editor
This story was originally posted by EBN: http://bit.ly/PHNBLY.

SIMPLIFIED.

SUPPLY CHAIN SOLUTIONS FROM AVNET

Your job is to develop and market innovative products; our job is to reduce your risk along the way.
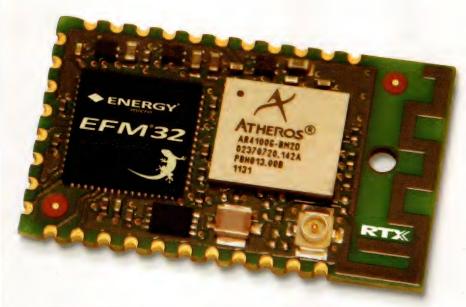
Even leading edge products need a cost-effective and dependable supply chain to be a success. That's where Avnet can help. Around the world and around the clock, you can rely on our expertise at every point in the supply chain – from supply chain assessment to solution design to global planning. Avnet's supply chain solutions save you inventory and infrastructure. That's money in the bank, and the edge you need to stay ahead of the competition. Simple, really.

What can we do for you? www.avnetexpress.com

Accelerating Your Success!™
1 800 332 8638 | www.avnetexpress.com | ▶ ⓔ @avnetdesignwire

AVNET®
electronics marketing

productroundup

RF AND WIRELESS

RTX SDK speeds time to market for ultralow-power Wi-Fi modules

RTX A/S has launched a software development kit for its latest ultralow-power Wi-Fi module, targeting M2M application developers. The SDK—comprising an API, demonstration applications, and an embedded firmware development suite with download tools—targets the RTX4100, a fully self-contained Wi-Fi module that RTX has released in cooperation with Energy Micro and Qualcomm Atheros. The RTX4100 small-form-factor, single-stream, 802.11b/g/n module has an onboard, low-power application processor. The software platform is available at no charge to end customers; pricing is $399 for the full RTX4100 development kit, $279 for an RTX4100 evaluation kit, and $99 for an RTX4100 wireless application board.

RTX A/S, www.rtx.dk

Cable-based antennas from Gore improve signal access for in-flight wireless

Cable-based antennas from WL Gore & Associates improve signal propagation without increasing hardware for in-flight applications. Easily installed along the length of the cabin ceiling, the antennas ensure reliable access. Signals are sent and received in frequencies ranging from 400 MHz to 6 GHz for compatibility with such standards as Bluetooth, DECT, DECT2, Global Star, GSM, Iridium, MMS, PDC, Tetra, UMTS, WLAN 802.11 a/b/g, and WiMAX.

WL Gore & Associates, www.gore.com

Near-field counters from Kaltman expedite device-frequency ID

Two additions to Kaltman Creations' Invisible Waves product line, the RF-id Station and the RD-id Solo, key in on wireless transmissions as near-field frequency counters. The RF-id Station self-contained case lets professional wireless-audio users stage and manage wireless microphones for live events. The rugged plastic carrying case has eight designated foam-formed slots for wireless microphones or belt-pack transmitters. Each RF-shielded slot has a backlit, near-field frequency-counter readout and a barograph signal-level display. With the RF-id Solo handheld, users can instantly confirm exact frequencies and transmission-signal strength. The unit works with digital and analog single-carrier transmissions in the 50-MHz to 2.5-GHz range.

Kaltman Creations, www.kaltmancreationsllc.com

TI demonstrates ZigBee Light Link on CC2530

Texas Instruments demonstrated the ZigBee Light Link standard on its CC2530 SOC for the first time at Computex Taipei in June. ZigBee Light Link standardizes wireless networked LED lighting systems for easy installation, management, and operation. The demo included a complete LED light-management system using a smartphone, low-cost gateway, and LED devices. An Android smartphone with a microSD

productroundup

robust wireless network nodes. Available now, the SOC is packaged in a ROHS-compliant, 6×6-mm QFN-40. The Z-Stack protocol stack for ZigBee Light Link is sampling; full availability is expected in the third quarter.
Texas Instruments, www.ti.com

card that integrated the CC2530 accomplished direct lighting control. The CC2530, a certified "golden unit" for the new ZigBee standard, lets designers and manufacturers rapidly build and test ZigBee Light Link products for interoperability. Integrating an RF transceiver, microcontroller, in-system programmable flash, and RAM, the CC2530 enables a low total bill-of-materials cost for

Atmel fields ref design for ZigBee Light Link

Atmel Corp's evaluation kit for ZigBee Light Link lets engineers manage multiple devices, such as lamps, switches, dimmers, timers, and remote controls, on one network. The reference

design integrates ATmega128RFA1 wireless microcontrollers on an RF4CE-EK evaluation kit with the ZigBee Light Link Profile software stack ported to the system. Engineers can also use the ZigBee Light Link reference implementation with Atmel's LED drivers for color mixing, system management, and advanced dimming in solid-state lighting apps. The software stack is available free; the RF4CE-EK sells for $479.
Atmel, www.atmel.com

Premo RFID antennas provide protection in automotive apps

Premo Group's TP0702U and TP0702UCAP SMD antennas for low-frequency, 20- to 150-kHz receiver applications provide 50-mV/A pk-pk/m sensitivity (for 7.2 mH). The copolyamide polyhexamethylene poly-
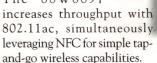
mer walls of the TP0702UCAP, gamma-radiated for high thermal stability and mechanical resistance, provide upper-side and lateral protection. The antenna has a NiZn ferrite core with high surface resistivity, providing stable behavior over a temperature range of −40° to +125°C.
Premo Group, www.grupopremo.com

Marvell 802.11ac mobile combo radio has MIMO, NFC

The Avastar 88W8897 low-power 802.11ac combination radio chip is the industry's first 802.11ac 2×2 combination radio chip, according to Marvell Technology Group. The chip delivers wireless, NFC (near-field communications), and Bluetooth 4.0, with MIMO, transmit beamforming, and support for Wi-Fi-certified Miracast. The 88W8897 increases throughput with 802.11ac, simultaneously leveraging NFC for simple tap-and-go wireless capabilities.
Marvell Technology Group, www.marvell.com

Ramtron dev kit revs wireless-memory capture/write speed

The MaxReader Development Kit lets engineers quickly evaluate, test, and prototype Ramtron International's MaxArias wireless memory for RF-enabled applications. The kit comprises the MaxArias WM72016 16-kbit wireless FRAM on a small-footprint, EPC (Electronic Product Code) Gen2 RF reader board, along with wireless-memory transponders and other hardware. MaxArias wireless FRAMs can capture block writes of up to 127 words without compromising data integrity or read/write speed and distance. The kit is available through Digi-

Key and Mouser Electronics in five regionally tuned frequency bands; the WM72016-based wireless-memory transponder board, WM72016-6-EVAL-915-ANT, is also available from the distributors as a stand-alone offering for a $25 suggested resale price.

Ramtron International,
www.ramtron.com

Maxim scalable QAM chip set delivers next-gen triple play

A high-density, wideband, scalable QAM (quadrature amplitude modulation) modular chip set comprising the MAX5880 digital upconverter and the MAX5882 RF digital-to-analog converter addresses the needs of the next-generation connected home. Delivering up to 32 times more broadband-cable capacity per RF port than conventional analog upconverters, the solution saves up to 93% in power dissipation per QAM channel. The chip set synthesizes up to 128 QAM channels across the 50-MHz to 1-GHz downstream cable spectrum. The MAX5880 is software-upgradable for capacity expansion. The chip set, in two 17×17-mm packages, and a high-speed data-converter evaluation platform and kit are available now.

Maxim Integrated Products,
www.maxim-ic.com

Freescale transistors target base stations

Freescale Semiconductor's Airfast transistors boost the efficiency, peak power, and signal bandwidth of next-generation small- and macro-cell base stations. The transistor portfolio now includes the AFT09S282N, the first Airfast 900-MHz, 28V LDMOS product; the AFT18S230S, a 28V transistor for symmetric or asymmetric Doherty applications, delivering 45% efficiency at 8-dB OBO (output backoff) and 17 dB of gain; the AFT21S230S, a 28V device in NI780S-6 packaging for use in symmetric or asymmetric Doherty applications; the AFT18HW355S, operating at 1805 to 1880 MHz or 1930 to 1995 MHz and enabling full-band, multicarrier operation; and the MMDS25254H ADAM (advanced Doherty alignment module), enabling optimization of Doherty amplifiers.

Freescale Semiconductor,
www.freescale.com

TriQuint 802.11ac Wi-Fi RF module enables fast download

The TQP6M9017, which TriQuint Semiconductor Inc says is the first 802.11ac Wi-Fi RF module for next-generation mobile devices, taps output-power technology improvements to allow a nearly 60% greater connectivity range. The highly integrated, dual-band wireless local-area network module provides a complete solution for 802.11 a/b/g/n/ac Wi-Fi and Bluetooth applications, simplifying RF design. Its 4×4-mm package integrates two power amplifiers for the 2.4- and 5-GHz frequency bands with a switch, filtering, baluns, and other components.

TriQuint Semiconductor,
www.triquint.com

What's your (single) point, youngster?

I was asked to help a young colleague with a product that was close to production. The circuit had a detector that connected to the virtual ground of an op-amp setup to convert detector current to voltage. The detector impedance was only about 1 kΩ; the feedback resistor's value was 100 MΩ. So, even though the op-amp circuit functioned as a transimpedance amplifier, it was also a voltage amplifier with a gain of 100,000. The unit was supposed to work with either batteries or a battery eliminator. We had tested it over the full input-voltage range, from an almost-dead battery to the highest possible eliminator voltage, but we didn't try a battery eliminator until late in the project. When we connected the eliminator, we saw a shift of approximately 100 mV.

Many theories existed regarding the cause. One engineer blamed it on breakdown phenomena involving the conductive coating on the inside of the case; others suggested strange current-leakage paths involving the same coating. Another posited RF rectification and was sure a few ferrite beads and some "good" capacitors would solve the problem.

I worked backward from the error and calculated that it would require only a 1-µV difference in the ground voltage between the op amp and the detector. I asked my younger colleague involved in the design whether he had used a true single-point ground, and he assured me that he had, sending me a copy of the artwork to prove it. I came back to him with the artwork and asked where the ADC, the detector, and the op amp's noninverting input connected to ground; he said each connected in a different place.

"That's not exactly a single-point ground," I said.

Rolling his eyes, the youngster said he'd done better than a single-point ground: He'd implemented a solid-copper ground plane for the analog section and a separate ground plane for the digital section. The common point was at the ADC. The whole analog ground plane carried only 100 µA and could not possibly develop 1 µV. The analog ground plane was, in fact, equivalent to a single-point ground.

"Have you got a unit that has the problem that I can play around with?" I asked.

"I've got about a hundred, and I've got to get them fixed. I'm looking into eliminating the error in software," he said. "You can have one, but please don't break it."

"Show me how to get the effect," I said, fetching a 4.5-digit digital voltmeter that was nearby and putting it on the 100-mV scale.

The youngster smirked. "You'll never see anything with that. It has only 10-µV resolution."
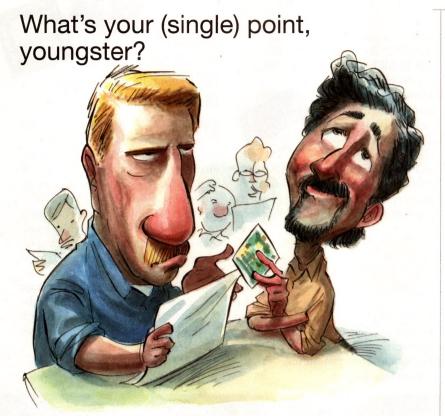
It was worth enduring all of his eye rolling when I found two points on the supposedly single-point analog ground that differed by 30 µV; this went away when I disconnected the battery eliminator.

The fix was simple. The ADC had differential inputs, so we tied its inverting input to the op amp's noninverting input. We moved the detector's ground connection to the op amp's noninverting input. The unit behaved flawlessly, even though its ground wasn't ground.

It turns out the youngster had inadvertently connected the return for the battery eliminator to the analog ground plane instead of the digital ground plane. When the eliminator was powering the unit, the entire load current flowed through the analog ground plane, causing a difference of approximately 1 µV between the old detector ground and the op amp's ground.

Two effects had ganged up on my colleague: a measly 1 µV of ground voltage and a low detector impedance that turned his transimpedance amplifier into a high-gain voltage amp.**EDN**

Roy McCammon is a senior technical specialist engineer with 3M's Track and Trace Division (St Paul, MN).